



# Finding Safety Violations of AI-Enabled Control Systems through the Lens of Synthesized Proxy Programs

JIEKE SHI, ZHOU YANG\*, and JUNDA HE, Singapore Management University, Singapore  
BOWEN XU, North Carolina State University, United States  
DONGSUN KIM, Korea University, South Korea  
DONGGYUN HAN, Royal Holloway, University of London, United Kingdom  
DAVID LO, Singapore Management University, Singapore

Given the increasing adoption of modern AI-enabled control systems, ensuring their safety and reliability has become a critical task in software testing. One prevalent approach to testing control systems is falsification, which aims to find an input signal that causes the control system to violate a formal safety specification using optimization algorithms. However, applying falsification to AI-enabled control systems poses two significant challenges: (1) it requires the system to execute numerous candidate test inputs, which can be time-consuming, particularly for systems with AI models that have many parameters, and (2) multiple safety requirements are typically defined as a conjunctive specification, which is difficult for existing falsification approaches to comprehensively cover.

This paper introduces SYNTHIFY, a falsification framework tailored for AI-enabled control systems, i.e., control systems equipped with AI controllers. Our approach performs falsification in a two-phase process. At the start, SYNTHIFY synthesizes a program that implements one or a few linear controllers to serve as a proxy for the AI controller. This proxy program mimics the AI controller's functionality but is computationally more efficient. Then, SYNTHIFY employs the  $\epsilon$ -greedy strategy to sample a promising sub-specification from the conjunctive safety specification. It then uses a Simulated Annealing-based falsification algorithm to find violations of the sampled sub-specification for the control system. To evaluate SYNTHIFY, we compare it to PSY-TALiRo, a state-of-the-art and industrial-strength falsification tool, on 8 publicly available control systems. On average, SYNTHIFY achieves a 83.5% higher success rate in falsification compared to PSY-TALiRo with the same budget of falsification trials. Additionally, our method is 12.8 $\times$  faster in finding a single safety violation than the baseline. The safety violations found by SYNTHIFY are also more diverse than those found by PSY-TALiRo, covering 137.7% more sub-specifications.

CCS Concepts: • **Software and its engineering** → **Software testing and debugging**; **Search-based software engineering**; • **Computing methodologies** → **Control methods**.

Additional Key Words and Phrases: Falsification, Search-based Testing, AI-enabled Control Systems, Program Synthesis

## 1 INTRODUCTION

Control systems, which utilize software controllers to maintain desired behaviors of particular mechanical systems, have been broadly deployed into diverse real-world applications. These systems have become an essential part of

\*Corresponding author.

---

Authors' Contact Information: Jieke Shi, jiekeshi@smu.edu.sg; Zhou Yang, zyang@smu.edu.sg; Junda He, jundahe@smu.edu.sg, Singapore Management University, Singapore; Bowen Xu, bxu22@ncsu.edu, North Carolina State University, United States; Dongsun Kim, darksw@korea.ac.kr, Korea University, South Korea; DongGyun Han, DongGyun.Han@rhul.ac.uk, Royal Holloway, University of London, United Kingdom; David Lo, davidlo@smu.edu.sg, Singapore Management University, Singapore.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2025 Copyright held by the owner/author(s).

ACM 1557-7392/2025/1-ART

<https://doi.org/10.1145/3715105>

everyday life and industrial production [114]. With the rapid advances of Artificial Intelligence (AI) in recent decades, there has been a huge demand for enhancing or even replacing traditional mathematical controllers (e.g., Proportional-Integral-Derivative [47]) with AI controllers (e.g., Deep Neural Networks) to achieve more optimized performance. We refer to control systems employing AI controllers as AI-enabled control systems. With their immense potential, AI-enabled control systems have inevitably been or will soon be adopted in various safety-critical domains, such as aerospace [14, 46], autonomous driving [53, 70], healthcare monitoring [50, 54], and even nuclear fusion [19]. Therefore, ensuring the reliability and safety of AI-enabled control systems in such contexts is of paramount importance.

To date, falsification has been shown to be an effective approach for testing complex control systems [18, 68, 93, 116, 125, 131]. Concretely, given a control system with a safety specification, usually written in Signal Temporal Logic (STL) [63], falsification aims to find an input signal that causes the control system to violate the specification. Falsification often uses an optimization algorithm to search for input signals with the goal of minimizing the quantitative robustness semantics of temporal logics [23]. These semantics quantify the degree to which an input signal deviates from violating the safety specification as a real number. A smaller number means that the corresponding input signal is more likely to violate a given safety specification, and a negative number suggests that the safety specification is violated. Falsification iteratively generates candidate input signals and executes them on the control system until an input with negative robustness semantics is found, revealing a violation of the safety specification. Various falsification methods have been proposed and adopted in practice [27, 52, 102], showing that falsification is effective in revealing specification violations in complex control systems.

However, applying falsification in AI-enabled control systems poses two challenges. First, falsification of AI-enabled control systems suffers from the scalability issue. Falsification requires the system under test to execute numerous candidate inputs until a safety violation is found. However, AI controllers are typically implemented as deep neural networks (DNNs) with numerous parameters, which are computationally expensive to execute. For example, DNN-based controllers in [134] have 4–50× longer execution time than classic linear controllers. Consequently, falsification can take a large amount of time to the extent that it becomes impractical.

Second, falsification of AI-enabled control systems fails to comprehensively cover conjunctive specifications. AI-enabled control systems emerge to handle highly complex control tasks, often with multiple safety specifications. For instance, a self-driving control system in [134] is required to maintain a car’s heading angle within  $90^\circ$  ( $\varphi_1$ ) while keeping the distance between the car and the centerline of the road below 2.0 meters ( $\varphi_2$ ). A common practice for falsifying such multiple specifications simultaneously is to combine them as a conjunctive specification ( $\varphi_1 \wedge \varphi_2$ ) [59]. However, directly falsifying such a conjunctive specification often produces less diverse safety violations. If one sub-specification is easily violated, the falsification process tends to repeatedly exploit that sub-specification to report violations and overlook others. As Sun et al. [95] discussed, it is desirable to find diverse violations that can cover as many sub-specifications as possible. Nevertheless, the existing state-of-the-art falsification tools like PSY-TALiRo [98] are not inherently designed to meet this challenge.

Inspired by recent advances in controller synthesis [64, 134], our key insight is that a complex AI controller can be approximated into a program that implements one or a few linear controllers, which is less computationally expensive to run. Of course, such a program with linear controllers may not precisely mirror the AI controller’s performance, but it does achieve similar functionality to some extent [64]. For example, experiments by Rajeswaran et al. [79] show that robots controlled by linear controllers may not walk like AI-enabled ones, but they can still stand up effectively. As analyzed by Zhu et al. [134], a program having similar functionality to an AI controller can also express similar safety properties. Thus, we propose to mitigate the scalability issue in the falsification of AI-enabled control systems by synthesizing computationally cheap programs as proxies for AI controllers.

Additionally, we consider that recent studies have highlighted the importance of the exploration-exploitation trade-off in falsifying conjunctive specifications [66, 80]. Specifically, falsification can either exploit a sub-specification that is easily violated to ensure successful falsification, or explore a sub-specification that has not

yet been falsified to boost the coverage of all sub-specifications within the conjunctive specification. The low coverage of sub-specifications is often a result of inadequate exploration and over-exploitation. Thus, we propose a strategy that balances the exploration and exploitation of sub-specifications, allowing falsification to more comprehensively cover the conjunctive specification.

Pursuant to the above insights, we propose a novel falsification framework named SYNTHIFY (SYNthesize-THEN-falsIFY). SYNTHIFY is tailored for AI-enabled control systems using Deep Reinforcement Learning (DRL) algorithms, which are considered as the state-of-the-art in various control tasks. To facilitate efficient and effective falsification of AI-enabled control systems, SYNTHIFY combines two distinct techniques: (1) sketch-based program synthesis [91] to effectively synthesize a proxy program from the DRL-based AI controller. The proxy program implements one or a few linear controllers to mimic the functionality of the AI controller but is much cheaper to run; (2)  $\epsilon$ -greedy strategy [55], a classic strategy for balancing exploration and exploitation in the decision-making process. This strategy helps SYNTHIFY sample a promising sub-specification from the given conjunctive safety specification. SYNTHIFY then utilizes the synthesized program as a proxy for the AI controller in the control system and employs a Simulated Annealing-based algorithm to guide the search for violations of the sampled sub-specification.

We evaluate SYNTHIFY on 8 publicly available control systems from the literature [134] and compare it to a state-of-the-art falsification tool PSY-TALiRo [98], which recently won the 2022 ARCH-COMP prize for its technical achievements and excellent performance in falsification [30]. Our results show that, given the same budget of falsification trials, SYNTHIFY achieves a 83.5% higher falsification success rate than PSY-TALiRo, while also being 5.6 $\times$  faster than the baseline at completing the same number of falsification trials. Given the same time budget, SYNTHIFY reveals 7.8 $\times$  more violations of safety specifications than PSY-TALiRo. Moreover, our method is 12.8 $\times$  faster than the baseline at finding one single violation. The safety violations found by SYNTHIFY are also more diverse than those found by PSY-TALiRo, improving 137.7% the coverage of the conjunctive specification on average. These results demonstrate the effectiveness and efficiency of the proposed method in testing the safety of AI-enabled control systems.

The contributions of this paper are summarized as follows:

- **Insight:** We are the first to combine sketch-based program synthesis and  $\epsilon$ -greedy strategy to simultaneously address the scalability and sub-specification coverage issues in the falsification of AI-enabled control systems.
- **Technique:** We propose and implement SYNTHIFY, an innovative falsification framework that can both efficiently and effectively uncover safety violations of AI-enabled control systems.
- **Evaluation:** We evaluate SYNTHIFY on 8 control systems and compare it to PSY-TALiRo, a state-of-the-art falsification tool. The results show that our method outperforms the baseline by a large margin in terms of both effectiveness and efficiency.

The remainder of the paper is organized as follows. Section 2 covers the preliminary information of our work. Section 3 shows an example to motivate our work, while Section 4 describes in detail how SYNTHIFY works. Section 5 presents the experimental setup and research questions. Section 6 analyzes the experimental results. Section 7 presents a number of additional analyses and discussions. Section 8 presents related work, and Section 9 concludes our paper and provides information about our replication package.

## 2 PRELIMINARIES

This section provides an introduction to AI-enabled control systems. It also gives a brief description of Signal Temporal Logic (STL) as a formal specification language for control systems, and falsification, a common and important testing method used in the safety analysis of control systems.

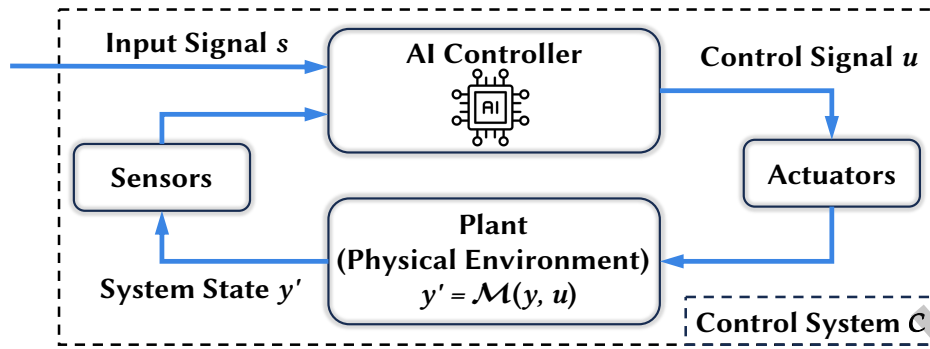


Fig. 1. Architecture of an AI-enabled control system.

## 2.1 AI-enabled Control System

Control systems utilize software controllers to regulate the behavior of mechanical systems based on real-time feedback from their physical environments [124]. Traditionally, controllers were designed with classic algorithms like Proportional-Integral-Derivative (PID) [47]. However, with the growing interest in AI, there has been an exploration of AI controllers powered by deep neural networks. We refer to control systems equipped with AI controllers as AI-enabled control systems.

Figure 1 shows a typical architecture of AI-enabled control systems consisting of a plant, a controller, sensors, and actuators. The plant  $\mathcal{M}$  is a physical environment whose next state  $y'$  is determined by the current state  $y$  and the control signal  $u$ . The control signal  $u$  is issued by an AI controller and applied to the plant by actuators. Integrating AI controllers allows control systems to adapt to more complex situations and often achieve superior results compared to traditional methods [43, 60]. However, this integration also introduces new risks, such as the potential for system failures that could violate safety requirements. These risks underscore the need for our work in analyzing the safety and reliability of AI-enabled control systems.

The AI controller is typically implemented as a deep neural network, and Deep Reinforcement Learning (DRL) is currently considered the state-of-the-art method for training it [26, 60, 93]. Our study focuses on control systems with DRL-based AI controllers. DRL is a subfield of machine learning that focuses on learning optimal decision strategies through interactions with the environment [111]. In DRL, an agent learns the best policy by observing states and selecting the best output to maximize the cumulative reward. Concretely, considering a DRL-based AI controller as an agent, it receives the current state of the plant and outputs a control signal to transition the system to a new state. A policy, a deep neural network taking the system's state as input, determines the control signal that the agent outputs. The reward for each output under a state depends on the feedback from the plant. Through repeated interactions with the plant, the agent is trained to maximize the reward it receives, learning and adapting its policy to make better decisions over time. After training, the agent with its best policy is deployed as an AI controller in the control system. DRL-based AI controllers have found applications in many critical areas such as autonomous driving [53, 113], robotics control [41, 48], etc.

## 2.2 Signal Temporal Logic

The safety requirements of a control system are usually specified using Signal Temporal Logic (STL) [63], a widely-adopted formal specification language. STL describes how a control system should behave by combining atomic propositions and formulas using logical and temporal operators. Formally, the basic syntax of an STL

formula  $\varphi$  is defined as:

$$\varphi ::= \mu \mid \perp \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \mathcal{U}_{\mathcal{I}} \varphi_2 \quad (1)$$

where  $\mu ::= f(s) \sim c$  is an atomic proposition,  $f(s)$  is a real-valued function over the variables  $s$ ,  $\sim \in \{<, \leq, =, \geq, >\}$  is a comparison operator, and  $c$  is a constant. Additionally,  $\perp$  represents “false”, while  $\neg$  and  $\wedge$  are the negation and conjunction operators, respectively.  $\mathcal{U}_{\mathcal{I}}$  is the “until” operator with a time interval  $\mathcal{I}$ , which is a closed non-singular interval in  $\mathbb{R}_{\geq 0}$ , i.e.,  $\mathcal{I} = [a, b]$  or  $[a, \infty]$  where  $a, b \in \mathbb{R}$  and  $a < b$ . From the basic definition, we can derive other logical and temporal operators, including  $\top \equiv \neg\perp$  (true),  $\vee \equiv \neg(\neg\varphi_1 \wedge \neg\varphi_2)$  (disjunction),  $\rightarrow \equiv \neg\varphi_1 \vee \varphi_2$  (implication),  $\diamond_{\mathcal{I}}\varphi \equiv \top \mathcal{U}_{\mathcal{I}} \varphi$  (eventually), and  $\square_{\mathcal{I}}\varphi \equiv \neg\diamond_{\mathcal{I}}\neg\varphi$  (always).

STL has quantitative robustness semantics [23] that map an input signal  $s$  and an STL specification  $\varphi$  to a real number, indicating the extent to which  $s$  satisfies  $\varphi$ . Formally, with respect to a signal  $s = (s_1, s_2, \dots, s_n)$ , the satisfaction of an atomic proposition  $\mu = f(s) \sim c$  at time  $t$  is defined as:

$$\begin{aligned} (s, t) \models \mu &\Leftrightarrow f(s_1[t], \dots, s_n[t]) \sim c \\ (s, t) \models \varphi_1 \wedge \varphi_2 &\Leftrightarrow (s, t) \models \varphi_1 \wedge (s, t) \models \varphi_2 \\ (s, t) \models \neg\varphi &\Leftrightarrow \neg((s, t) \models \varphi) \\ (s, t) \models \varphi \mathcal{U}_{\mathcal{I}=[a,b]} \varphi &\Leftrightarrow \exists t' \in [t+a, t+b] \text{ such that } (s, t') \models \varphi \wedge \forall t'' \in [t, t'], (s, t'') \models \varphi \\ (s, t) \models \diamond_{\mathcal{I}=[a,b]} \varphi &\Leftrightarrow \exists t' \in [t+a, t+b] \text{ such that } (s, t') \models \varphi \\ (s, t) \models \square_{\mathcal{I}=[a,b]} \varphi &\Leftrightarrow \forall t' \in [t, t+b] \text{ such that } (s, t') \models \varphi \end{aligned} \quad (2)$$

The larger the semantic value, the more robustly the input signal  $s$  satisfies the STL formula  $\varphi$ . A negative value indicates that  $s$  violates  $\varphi$ , with the degree of violation equal to the absolute value of the number. The intuition behind the robustness semantics is to measure the distances of the points on the system trace from the boundaries of the sets defined by the formula predicates using maximum or minimum functions. For example, let us consider the STL formula  $\square_{[0,1]}(a < 1.0)$ , specifying that a variable  $a$  must remain below 1.0 for the duration of the time interval  $[0, 1]$ . Suppose we have a system trace  $[(0.5, 0), (1.2, 1)]$ , where each element  $(a, \text{time})$  corresponds to  $a$ 's value and its associated time. For the first point, the value of  $a$  is below the threshold 1.0, so the distance to the boundary is simply the difference between the threshold and the value:  $1.0 - 0.5 = 0.5$ . For the second point,  $a$  goes up to the threshold, so the distance to the boundary is  $1.0 - 1.2 = -0.2$ . To compute the robustness semantics of the trace, we take the minimum distance across all points within the time interval, which in this case is a negative value of  $-0.2$ , indicating that the signal  $s$  violates the STL formula by 0.2. We refer interested readers to [23] for the complete definition of STL robustness semantics.

### 2.3 Falsification

Falsification is a well-established testing method used in the safety analysis of control systems. It aims to find inputs to a control system that cause the system to violate safety specifications expressed in STL. Concretely, given a control system  $C$  and an STL specification  $\varphi$ , falsification first searches for an input signal  $s$  and executes the system to obtain a system state trace. The trace is then used to compute the quantitative robustness semantics with respect to  $\varphi$ . If negative robustness is obtained, indicating a violation of  $\varphi$ , the falsification is successful and the input signal  $s$  is considered a falsifying input to the control system  $C$ . The search continues until a violation is found or a given budget (time or number of executions) is exhausted. Notably, falsification is typically performed in a black-box manner, i.e., the internal structures of the plant and the AI controller are unknown throughout the falsification process; only the input and the corresponding output signals can be observed.

Falsification is typically supported by optimization algorithms such as Simulated Annealing (SA) [1] shown in Algorithm 1. The algorithm starts with a randomly-initialized input signal  $s$  (line 2). It then iteratively perturbs it to generate a new input  $s'$  (line 6), which is then used to execute the system  $C$  to check if it is a falsifying input (lines 7-8). If  $s'$  does not cause the system to violate  $\varphi$ , i.e., obtaining non-negative robust semantics, the algorithm accepts  $s'$  into the next iteration with a probability that is determined by the temperature  $T$  (lines 9-11).



**Algorithm 1:** Simulated Annealing-based Falsification

---

**Require:** Control system  $C$ , STL specification  $\varphi$ , Input ranges  $\mathcal{S}$   
**Parameters:** Budget  $\mathcal{B}$ , Maximum number of iterations  $N$ , Cooling rate  $\alpha$ , and Initial temperature  $T$

```

1 repeat
2    $s \leftarrow \text{RANDOMSELECT}(\mathcal{S})$  /* Randomly initialize  $s$  from the given input ranges  $\mathcal{S}$  */
3    $rb \leftarrow \text{ROBUSTSEMANTICS}(C(s), \varphi)$ 
4   if  $rb < 0$  then return  $s$  /* Return  $s$  if getting a negative robustness */
5   for  $i \leftarrow 1$  to  $N$  do /* Simulated Annealing with  $N$  iterations */
6      $s' \leftarrow \text{PERTURB}(s)$ 
7      $rb_{s'} \leftarrow \text{ROBUSTSEMANTICS}(C(s'), \varphi)$ 
8     if  $rb_{s'} < 0$  then return  $s'$  /* Return  $s'$  if getting a negative robustness */
9      $\Delta \leftarrow rb_{s'} - rb, P \leftarrow e^{-\Delta/T}$  /* Accept  $s'$  if it is better or by probability */
10    if  $\Delta < 0$  or  $\text{RANDOM}(0, 1) < P$  then
11       $s \leftarrow s', rb \leftarrow rb_{s'}$ 
12     $T \leftarrow \alpha * T$  /* Update temperature */
13 until BUDGET  $\mathcal{B}$  (TIME OR NUMBER OF TRIALS) IS EXHAUSTED
14 return NONE /* No input that causes violations found */

```

---

**Algorithm 2:**  $\epsilon$ -greedy Strategy

---

**Require:** Options  $O = \{o_1, o_2, \dots, o_n\}$ , Expected rewards  $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$   
**Parameters:** Probability threshold  $\epsilon$

```

1 if  $\text{RANDOM}(0, 1) < \epsilon$  then
2    $o \leftarrow \text{RANDOMSAMPLE}(O)$  /* Randomly select an option */
3 else
4    $o \leftarrow \arg \max_{o_i \in O} \mathcal{R}_i, i \in [1, n]$  /* Select the option with the highest expected reward */
5 return  $o$  /* Return the selected option */

```

---

The algorithm terminates when the budget  $\mathcal{B}$  is exhausted or when a falsifying input signal  $s$  is found, witnessing the violation of  $\varphi$ .

## 2.4 $\epsilon$ -greedy Strategy

When iteratively selecting among multiple options, a key challenge is to balance exploration (i.e., trying new options that might yield better results) and exploitation (i.e., choosing the best option based on current knowledge). In this context, the  $\epsilon$ -greedy strategy [96] offers a straightforward yet effective solution. The  $\epsilon$  in its name represents a probability that governs the trade-off between exploration and exploitation. Formally, consider a set of  $n$  options  $O = \{o_1, o_2, \dots, o_n\}$ , where the performance of each option is evaluated through reward feedback, we can obtain the knowledge of the expected reward  $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$  of each option by averaging the rewards received from multiple attempts so far. Then, as shown in Algorithm 2, with a probability of  $1 - \epsilon$ , the strategy exploits the current knowledge by selecting the option with the highest expected reward (line 4). Alternatively, with probability  $\epsilon$ , it explores by selecting an option at random (line 1), enabling the discovery of potentially better alternatives.

As shown by Zhang et al. [130], certain sub-specifications within a conjunctive specification are significantly easier to falsify than others as they are more sensitive to input signals. This sensitivity causes falsification




Fig. 2. An example of self-driving control system.

algorithms to disproportionately focus on these easier sub-specifications, reducing the likelihood of falsifying more difficult ones and potentially resulting in an incomplete safety analysis. Our approach treats each sub-specification as an independent option and incorporates the  $\epsilon$ -greedy strategy to encourage the falsification of sub-specifications that have not yet been falsified, thereby ensuring a more comprehensive safety analysis. A detailed explanation of how the  $\epsilon$ -greedy strategy works in our approach is provided in Section 4.2.

### 3 MOTIVATING EXAMPLE

This section uses the example of a self-driving system provided by Zhu et al. [134] to better illustrate the two main challenges associated with the falsification of AI-enabled control systems. We then show how SYNTHIFY effectively addresses these challenges.

**Self-Driving System.** Figure 2 presents a self-driving control system that uses a DRL-based AI controller to control a car to follow a given route. The system has a 2-dimensional state  $[\eta, d]$ , where  $\eta$  is the angle between the car's heading and the road's centerline, and  $d$  is the distance between the car's front and centerline. The AI controller takes the state as input and outputs a 1-dimensional control signal  $\delta$ , which is used to control the steering angle. A regulator ensures the signal  $\delta$  remains within the range of  $[-10, 10]$ , preventing rapid or extreme steering adjustments. As described by Zhu et al. [134], the car keeps moving at a constant speed, and the safety property of this system is to ensure that the car never veers into canals found on both sides of the road. Formally, the self-driving system is safe if the car's heading angle  $\eta$  remains within a threshold of  $90^\circ$  and the distance  $d$  between the car's front and the road's centerline stays below 2.0 meters. Such safety requirements are defined as the following conjunctive specification written in STL:

$$\varphi_{safety} \equiv \square_{[0,200]} (|\eta| < 90^\circ \wedge |d| < 2.0) \quad (3)$$

where the temporal operator  $\square_{[0,200]}$  indicates that the safety property must hold continuously for 200 time steps throughout the system's execution. This control system, though simple, serves as a representative example of many practical control systems that utilize AI models to govern the system's behavior, such as the Udacity self-driving car [103]. We employ this example to clearly illustrate specific challenges that motivate the design of SYNTHIFY, while more complex control systems are provided in Section 5.1.

**Key Challenges.** We use the state-of-the-art falsification tool, PSY-TALiRo [98], which employs the Simulated Annealing-based falsification algorithm outlined in Algorithm 1, to find safety violations of the self-driving system. We set the maximum number of falsification trials to 50 and feed the above-mentioned safety specification

Table 1. Running time and sub-specification violation coverage of the SOTA falsification tool PSY-TALiRo on the self-driving system over the 50 trials.

Running time			Sub-specification Coverage
AI Controller	Plant Execution	Algorithm	
<b>408.7s / 78.8%</b>	73.8s / 14.2%	36.1s / 7.0%	<b>50.0%</b>

into PSY-TALiRo. The results of running time and sub-specification violation coverage are presented in Table 1. Following Sun et al. [95], the sub-specification violation coverage is defined as follows:

**Definition 1: Sub-specification Violation Coverage.** Given a set of violations  $\mathcal{V}$  against a safety specification  $\varphi$  having  $n$  sub-specifications, the sub-specification violation coverage is the ratio of violated sub-specifications to the total number of sub-specifications, i.e.,

$$\text{Sub-specification Coverage} = \frac{|\{\varphi_i \mid \exists v \in \mathcal{V}, \neg(v \models \varphi_i), i \in [1, n]\}|}{n} \times 100\% \quad (4)$$

From Table 1, we can see that in the 50 falsification trials, the execution of the AI controller is the most time-consuming part, taking 408.7 seconds, which accounts for 78.8% of the total running time. This significant duration is attributed to the AI controller’s numerous parameters, resulting in high execution costs. This in turn makes the overall falsification time too long to be practical. Thus, the first challenge in falsification of AI-enabled control systems concerns the scalability issue, which can be addressed by reducing the cost of executing the AI controllers.

Additionally, in the example of the self-driving system, the safety requirements in Equation 3 is a conjunctive specification, which can be divided into four sub-specifications as follows:

$$\begin{aligned} \varphi_{safety} \equiv & \{\varphi_1 \equiv \square_{[0,200]} (\eta > -90^\circ), \varphi_2 \equiv \square_{[0,200]} (\eta < 90^\circ), \\ & \varphi_3 \equiv \square_{[0,200]} (d > -2.0), \varphi_4 \equiv \square_{[0,200]} (d < 2.0)\} \end{aligned} \quad (5)$$

We then calculated the sub-specification violation coverage of PSY-TALiRo, as shown in Table 1. Within 50 trials, the safety violations found by PSY-TALiRo could only falsify 50% of the sub-specifications within the conjunctive specification, specifically the first 2 of the 4 sub-specifications in Equation 5. The remaining 2 sub-specifications were missed by PSY-TALiRo. This result suggests that another challenge in falsification of AI-enabled control systems is the lack of comprehensiveness in covering all the sub-specifications required for ensuring safety.

**How & Why SYNTHIFY Works.** Our SYNTHIFY is purpose-built to address the above two challenges in the falsification of AI-enabled control systems. To improve scalability, we propose an Evolution Strategy-based synthesis algorithm (Section 4.1), which allows SYNTHIFY to synthesize a program to act as a proxy for the AI controller in the control system. To offer a clearer insight into the synthesized proxy program, we present an example in the self-driving scenario:

```

1 def self_driving( $\eta$ ,  $d$ ):
2      $\delta$  = 0.20706786 *  $\eta$  - 0.31286586 *  $d$  - 0.27174068
3     return  $\delta$ 

```

The proxy program includes a linear controller of the system states  $\eta$  and  $d$  and return a control signal  $\delta$  to control the car’s steering angle. It closely emulates the control behavior of the original AI controller, while being much more efficient to execute. During the falsification process, the proxy program replaces the original AI controller in the control system, and the AI controller is only used to verify if a real safety violation has been found. As a result, the execution cost of the AI controller is considerably reduced, effectively mitigating the scalability issue.



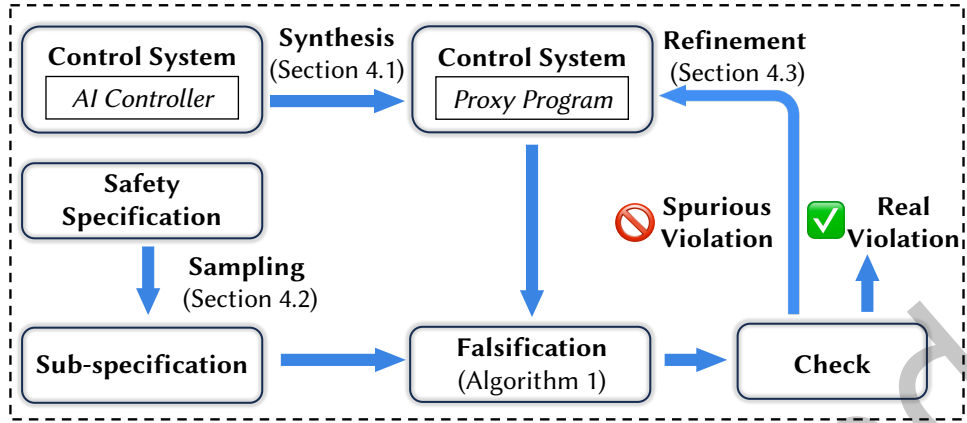


Fig. 3. Overview of SYNTHIFY.

Furthermore, to enhance the sub-specification violation coverage, SYNTHIFY proposes an  $\epsilon$ -greedy strategy-based method, which involves sampling a promising sub-specification from the conjunctive specification for the subsequent falsification process. The  $\epsilon$ -greedy strategy incorporates a probability of randomly selecting a sub-specification, which gives a chance for the falsification process to explore sub-specifications that have not been falsified yet, significantly improving the sub-specification coverage. More information on the  $\epsilon$ -greedy strategy can be found in Section 4.2.

## 4 DETAILED APPROACH

This section provides a detailed introduction to SYNTHIFY. Figure 3 shows an overview of our approach. Given a control system with an AI controller, the approach first approximates a proxy program from the controller to replace the controller (Section 4.1). Next, we separate the conjunctive safety specifications and sample one from them as the input to the subsequent falsification algorithm (Section 4.2). Finally, we employ the Simulated Annealing-based algorithm to search for violations of the sampled specification (Algorithm 1). If a violation is detected in the control system with the proxy program, we check whether it is spurious, i.e., it does not reveal any violations in the original control system. If spurious, we use it to refine the proxy program (Section 4.3); otherwise, we report it as a violation of safety specifications.

### 4.1 Synthesis

To facilitate the efficient falsification of AI-enabled control systems, SYNTHIFY first synthesizes a proxy program for the AI controller that is computationally more efficient. The proxy program replaces the AI controller in the control system during falsification. Our synthesis procedure consists of two steps: (1) constructing a sketch that can express the functionality of the AI controller; (2) using an Evolution Strategy-based algorithm to synthesize the blank coefficients in the sketch.

**4.1.1 Sketching.** Our approach models the synthesis of proxy programs for AI controllers as a sketch-based program synthesis problem [91]. A sketch is a partial code template with blank holes that are to be filled in. The synthesizer then searches for statements or parameters that can fill these holes. In our study, we create a sketch following [134] that allows us to implement one or a few linear controllers over a collection of variables

representing the system states of the control system. The sketch is parameterized by a matrix of coefficients  $\theta$  and is defined as follows:

```

1 def Proxy[ $\theta$ ]( $x_1, x_2, \dots, x_n$ ):
2    $y_1 = \theta_{[1,1]} * x_1 + \theta_{[1,2]} * x_2 + \dots + \theta_{[1,n]} * x_n + \theta_{[1,n+1]}$ 
3    $y_2 = \theta_{[2,1]} * x_1 + \theta_{[2,2]} * x_2 + \dots + \theta_{[2,n]} * x_n + \theta_{[2,n+1]}$ 
4   ...
5    $y_m = \theta_{[m,1]} * x_1 + \theta_{[m,2]} * x_2 + \dots + \theta_{[m,n]} * x_n + \theta_{[m,n+1]}$ 
6   return  $y_1, y_2, \dots, y_m$ 

```

where  $x_1, x_2, \dots, x_n$  are the system state variables in control systems,  $y_1, y_2, \dots, y_m$  are the outputs of the synthesized program.  $\theta$  is a matrix of coefficients with size  $m \times (n + 1)$ , where  $m$  and  $n$  are the number of outputs of the AI controller and the number of system state variables in the control system, respectively.  $\theta_{[i,j]}$  represents the  $j$ -th coefficient in the  $i$ -th equation. Note that the last coefficient  $\theta_{[* , n+1]}$  is a constant term that is not associated with any variable. Such a constant can be viewed as a bias term that can allow linear programs to express a wider range of functions [57]. This definition allows us to flexibly create program sketches for any control system by adjusting the number of equations and variables, as long as the system's inputs and outputs are known. We believe this can be done with low difficulty.

Although our sketch implementing linear controllers may resemble a simple neural network with weights and biases, there are key differences in their representations, with ours being simpler and more efficient. Specifically, our linear controller computes the output as a linear combination of the inputs, typically expressed as  $y = W \cdot x + b$ , where  $W$  is the coefficient matrix,  $x$  is the input vector, and  $b$  is the bias. In contrast, a neural network generally consists of multiple layers, each potentially incorporating nonlinear activation functions (e.g., tanh and sigmoid). Even a basic feedforward neural network with only one input layer and one hidden layer involves more complex computations than a linear controller, represented as  $y = \sigma(W_2 \cdot (\sigma(W_1 \cdot x + b_1) + b_2))$ , where  $W_1$  and  $W_2$  are the weight matrices,  $b_1$  and  $b_2$  are the bias vectors in the input and hidden layers, and  $\sigma$  is a nonlinear activation function. Neural networks with this structure often incur higher computational costs during prediction and potentially take more time to synthesize, so we opt for linear controllers in our sketch for greater simplicity and efficiency.

**4.1.2 Evolution Strategy.** Once we have constructed the sketch, we then employ an Evolution Strategy (ES)-based synthesis algorithm to synthesize the coefficients in the sketch. ES is a well-known optimization algorithm that has been successfully applied to solve program synthesis [90] and controller synthesis problems [82]. The ES-based synthesis algorithm is presented in Algorithm 3. The algorithm first randomly initializes the coefficients  $\theta$  in the sketch (line 1), then performs the following steps iteratively:

- It first generates a matrix of random noise  $N$  of size  $m \times |\theta|$ , from a standard normal distribution (line 3). When generating the matrix, we adopt mirror sampling [13, 64] to improve the performance of the algorithm. The algorithm first generates a random noise matrix of size  $\frac{m}{2} \times |\theta|$ , and then mirrors the matrix by negating each value to get another matrix with the same size. The algorithm then concatenates the two matrices to obtain a matrix of random noise of size  $m \times |\theta|$  (line 4);
- For each row in the random noise matrix  $N$ , the algorithm perturbs the current coefficients  $\theta$  by adding the row to the current coefficients  $\theta$  and obtains a new matrix of coefficients  $\theta'$  (line 7);
- The algorithm proceeds by obtaining the fitness matrix  $R$  for each perturbed coefficient  $\theta'$ . It executes the control system  $C$  using both the proxy program  $\mathcal{P}[\theta'](x_1, \dots, x_n)$  and the AI controller  $A^c$  for once. The fitness is then computed as the distance between the outputs of the proxy program and the AI controller (line 8);

**Algorithm 3:** Evolution Strategy-based Synthesis

---

**Require:** AI-enabled control system  $C$ , AI controller  $A^c$ , Proxy program sketch  $\mathcal{P}[\theta](x_1, \dots, x_n)$  with blank coefficients  $\theta$

**Parameters:** Maximum number of time steps  $t$ , Maximum number of iterations  $k$ , Population size  $m$ , Noise standard deviation  $\sigma$ , and Learning rate  $\alpha$

```

1  $\theta \leftarrow \text{RANDOMGENERATE}(|\theta|)$  /* Initialize coefficients */
2 repeat
3    $\mathbf{N} \leftarrow \text{RANDOMGENERATE}(m/2, |\theta|)$  /* Generate  $m \times |\theta|$  random noise by mirror sampling */
4    $\mathbf{N} \leftarrow \text{CONCATENATE}(\mathbf{N}, -\mathbf{N})$ 
5   for  $j \leftarrow 1$  to  $t$  do /* Update coefficients for  $t$  time steps */
6     for  $i \leftarrow 1$  to  $m$  do
7        $\theta' \leftarrow \theta + \mathbf{N}_i$  /* Perturb the current coefficients */
8        $\mathbf{R} \leftarrow \text{DISTANCE}(C, A^c, \mathcal{P}[\theta'](x_1, \dots, x_n))$  /* Calculate the fitness */
9        $\mathbf{R} \leftarrow (\mathbf{R} - \text{MEAN}(\mathbf{R})) / \text{STANDARDDEVIATION}(\mathbf{R})$  /* Standardizes the fitness */
10       $\theta \leftarrow \theta + \alpha \frac{1}{m\sigma} \sum_{i=1}^m \mathbf{N}_i \cdot \mathbf{R}_i$  /* Update coefficients */
11 until MAXIMUM NUMBER OF ITERATIONS  $k$  IS REACHED
12 return  $\mathcal{P}[\theta](x_1, \dots, x_n)$  /* Program with synthesized coefficients */

```

---

- The algorithm standardizes the fitness  $\mathbf{R}$  (line 9) and then updates the coefficients  $\theta$  by the following update function (line 10):

$$\theta \leftarrow \theta + \alpha \frac{1}{m\sigma} \sum_{i=1}^m \mathbf{N}_i \cdot \mathbf{R}_i \quad (6)$$

where  $\alpha$  is the learning rate,  $\sigma$  is the noise standard deviation, and  $\mathbf{N}_i$  and  $\mathbf{R}_i$  are the  $i$ -th row of the random noise matrix  $\mathbf{N}$  and the fitness  $\mathbf{R}$ , respectively. Such an update is to perform a stochastic gradient estimate step introduced in [82], which has been shown to be effective in optimizing coefficients and achieving better performance than gradient-based optimization algorithms [64, 82]. Unlike gradient-based methods, however, this algorithm does not require gradient information derived from partial derivatives, making it more efficient in certain cases.

The update procedure is executed for  $t$  time steps at each iteration out of the  $k$  iterations, i.e.,  $\theta$  is updated  $n \times k$  times. The algorithm either terminates when the maximum number of iterations  $k$  is reached or the fitness of the proxy program  $\mathcal{P}[\theta](x_1, \dots, x_n)$ , i.e., the distance between the outputs of the proxy program and the AI controller, is smaller than a predefined threshold, which indicates that the synthesized program is sufficiently similar to the AI controller. The algorithm returns the proxy program  $\mathcal{P}[\theta](x_1, \dots, x_n)$  with the synthesized coefficients  $\theta$ . All the hyperparameter settings in this algorithm can be found in Section 5.1.

Here one may wonder why we do not directly perform formal verification on the synthesized proxy program but instead perform falsification with it. This is because safety specifications are typically built on the plant rather than the controllers. The outputs of the AI controllers or proxy programs do not always reflect the safety properties of the system. Therefore, we apply falsification to analyze the control system to identify safety violations.

## 4.2 Sampling

To achieve more comprehensive coverage of the conjunctive safety specification during falsification, SYNTHIFY first separates the conjunctive specifications into sub-specifications. We can then select one sub-specification from the set of all sub-specifications to falsify. However, blindly exploring the sub-specification that has not been falsified to falsify is not effective, as some sub-specifications are more difficult to falsify than others or even

**Algorithm 4:**  $\epsilon$ -greedy Strategy-based Sub-specification Sampling

---

**Require:** Sub-specifications  $\Phi$ , Historical lowest robustness semantics for each sub-specification  $\mathcal{R}$   
**Parameters:** Probability threshold  $\epsilon$

```

1 if RANDOM(0, 1) <  $\epsilon$  then
2   |  $\varphi \leftarrow \text{RANDOMSAMPLE}(\Phi)$                                /* Sample one sub-specification randomly */
3 else
4   |  $i \leftarrow \arg \min_{i \in \{1, \dots, |\mathcal{R}|\}} \mathcal{R}$            /* Sample the sub-specification with the lowest robustness */
5   |  $\varphi \leftarrow \Phi_i$ 
6 return  $\varphi$                                                  /* Return a sub-specification for subsequent falsification */
```

---

impossible to falsify. Always forcing the falsification process to work on such sub-specifications can result in a waste of resources. Conversely, always exploiting previously falsified sub-specifications may not be effective either, as it may result in low sub-specification violation coverage. Thus, we propose an  $\epsilon$ -greedy Strategy-based Sampling method to intelligently decide which sub-specifications to falsify, balancing exploration and exploitation, thereby boosting the sub-specification coverage effectively.

The basic concept of the  $\epsilon$ -greedy strategy has been introduced in Section 2.4. Here, we describe its application in the context of falsification, as outlined in Algorithm 4. The primary adaptations of the  $\epsilon$ -greedy strategy to the sub-specification sampling process involve two key changes: (1) defining the reward as the historical robustness semantics of each sub-specification from prior falsification attempts—rather than using expected rewards, we track the lowest robustness value for each sub-specification; and (2) modifying the exploitation phase to select the sub-specification with the lowest robustness value instead of the highest reward, since a lower robustness value indicates a higher likelihood that the sub-specification may be violated in future attempts.

Given a set of sub-specifications  $\Phi$ , our algorithm begins by determining whether to explore by randomly selecting a sub-specification, using a randomly-generated number (line 1). If this number is less than the exploration probability  $\epsilon$ , the algorithm randomly selects a sub-specification from  $\Phi$  (line 2). Otherwise, it proceeds with exploitation, selecting the sub-specification with the lowest robustness value based on the historical robustness semantics  $\mathcal{R}$  (line 4). The selected sub-specification is then returned for subsequent falsification. The algorithm is called at each iteration of the falsification process to decide which sub-specification to falsify.

### 4.3 Falsification & Refinement

**4.3.1 Falsification.** We employ the Simulated Annealing (SA)-based falsification algorithm in SYNTHIFY, which is already described in Algorithm 1. Note that in SYNTHIFY the falsification algorithm is performed on the control system with a synthesized proxy program rather than the AI controller. Since the proxy program is not with the exact same functionality as the original AI controller, the SA-based falsification algorithm may find some spurious violations. Therefore, we use the original AI controller to check whether the violations are spurious or not. If the violations are spurious, we refine the proxy program to perform more similarly to the original AI controller. Otherwise, we terminate the falsification process and report the found violations.

**4.3.2 Refinement.** Once we have a spurious violation, we use the spurious falsifying input to refine the proxy program, i.e., slightly changing the coefficients in the proxy program to make its outputs more similar to those of the AI controller. The refinement procedure follows Algorithm 5. It still follows the ES-based synthesis presented in Algorithm 3, but introduces two major differences. First, instead of starting with a random matrix, the refinement begins with the current proxy program and its synthesized coefficients. Second, rather than initializing the control system randomly and collecting system states for fitness evaluation, the spurious falsifying input  $s$  is used to initialize the system, after which fitness evaluation occurs (lines 3-7). The rest of the procedure remains the same

**Algorithm 5:** Refinement of Synthesized Proxy Program

---

**Require:** AI-enabled control system  $C$ , AI controller  $A^c$ , Proxy program sketch  $\mathcal{P}[\theta](x_1, \dots, x_n)$  with synthesized coefficients  $\theta$ , Spurious falsifying input  $s$

**Parameters:** Maximum number of time steps  $t$ , Population size  $m$ , Noise standard deviation  $\sigma$ , and Learning rate  $\alpha$

```

1  $N \leftarrow \text{RANDOMSAMPLE}(m/2, |\theta|)$  /* Generate  $m \times |\theta|$  random noise by mirror sampling */
2  $N \leftarrow \text{CONCATENATE}(N, -N)$ 
3  $C \leftarrow C(s)$  /* Initialize the control system with the spurious falsifying input */
4 for  $j \leftarrow 1$  to  $t$  do /* Update coefficients for  $t$  time steps */
5   for  $i \leftarrow 1$  to  $m$  do
6      $\theta' \leftarrow \theta + N_i$  /* Perturb the current coefficients */
7      $R \leftarrow \text{DISTANCE}(C, A^c, \mathcal{P}[\theta'](x_1, \dots, x_n))$  /* Calculate the fitness */
8      $R \leftarrow (R - \text{MEAN}(R)) / \text{STANDARDDEVIATION}(R)$  /* Standardizes the fitness */
9      $\theta \leftarrow \theta + \alpha \frac{1}{m\sigma} \sum_{i=1}^m N_i \cdot R_i$  /* Update coefficients */
10 return  $\mathcal{P}[\theta](x_1, \dots, x_n)$  /* Program with refined coefficients */

```

---

as the ES-based synthesis procedure. Importantly, the refinement is applied only once after a spurious violation is detected during the falsification process to avoid overfitting the proxy program to the spurious input. Such a refinement procedure is inspired, at a high level, by the counterexample-guided abstraction refinement loop in formal verification [17, 68, 84]. The difference is that instead of using the counterexample to refine a model abstraction, we use the spurious falsifying input to refine the synthesized proxy program.

## 5 EVALUATION SETTINGS

Our evaluation aims to answer the following research questions:

**RQ1 (Effectiveness): How effective is SYNTHIFY in revealing violations of safety specifications?** Although SYNTHIFY is highlighted to resolve the issues of efficiency and sub-specification coverage in falsification of AI-enabled control systems, it is still unclear how effective SYNTHIFY is in finding safety violations. We evaluate the ability of SYNTHIFY in 8 AI-enabled control systems and compare it to our baseline PSY-TALiRo [98] which is introduced in Section 5.1.

**RQ2 (Efficiency): How efficient is SYNTHIFY in detecting violations of safety specifications?** One of the most innovative aspects of SYNTHIFY is its use of a synthesized proxy program to address the scalability issue in falsification of AI-enabled control systems. To demonstrate the advantage of this approach, we evaluate the time efficiency of SYNTHIFY by comparing it to the baseline PSY-TALiRo that does not incorporate program synthesis.

**RQ3 (Sub-specification Violation Coverage): How comprehensive is SYNTHIFY in covering safety specifications?** Another key feature of SYNTHIFY is its capability to cover many sub-specifications within conjunctive specifications. To demonstrate this advantage, we compare the sub-specification violation coverage of SYNTHIFY (i.e., how many sub-specifications in conjunctive safety specifications can be covered by its detected violations) with that of the baseline that does not employ sub-specification sampling.

### 5.1 Benchmark & Baseline

**5.1.1 Benchmark.** Table 2 shows 8 of the 15 control systems from Zhu et al. [134]. We adopt these systems in our evaluation and exclude the other 7 as they already show perfect performance, i.e., never violate any safe specification, according to the experiments in [134]. The 8 selected ones are representative of real-world control systems, such as Segway transporters [83], camera drones [21], and autonomous driving cars [103], and all are

Table 2. Control systems and their safety specifications. The detailed description of these systems and specifications are in Section 5.1. The table also shows the number of sub-specifications, inputs, outputs, and neural networks sizes of AI controllers.

Control System	Safety Specification	#Sub-Specifications	#Inputs/ #Outputs	AI Controller Size	
				#Layers	#Neurons
CartPole	$\square_{[0,200]} ( \eta  < 90^\circ \wedge  \delta  < 0.3 \wedge  v_1  < 0.3 \wedge  v_2  < 0.5)$	8	4/1	3	[300, 250, 200]
Pendulum	$\square_{[0,200]} ( \eta  < 90^\circ \wedge  \omega  < \pi/2)$	4	2/1	3	[280, 240, 200]
Quadcopter	$\square_{[0,300]} ( \eta_1  < \pi/2 \wedge  \eta_2  < \pi/2)$	4	2/1	2	[240, 200]
Self Driving	$\square_{[0,200]} ( \eta  < 90^\circ \wedge  d  < 2.0)$	4	2/1	3	[300, 250, 200]
Lane Keeping	$\square_{[0,300]} ( d  < 0.9)$	2	4/2	2	[240, 200]
4-Car Platoon	$\square_{[0,200]} ( a  \leq 2.0 \wedge  b  \leq 0.5 \wedge  c  \leq 0.35 \wedge  d  \leq 0.5 \wedge  e  \leq 1.0 \wedge  f  \leq 0.5 \wedge  g  \leq 1.0)$	14	7/4	3	[500, 400, 300]
8-Car Platoon	$\square_{[0,200]} ( a  \leq 2.0 \wedge  b  \leq 0.5 \wedge  c  \leq 1.0 \wedge  d  \leq 0.5 \wedge  e  \leq 1.0 \wedge  f  \leq 0.5 \wedge  g  \leq 1.0 \wedge  h  \leq 0.5 \wedge  i  \leq 1.0 \wedge  j  \leq 0.5 \wedge  k  \leq 1.0 \wedge  l  \leq 0.5 \wedge  m  \leq 1.0 \wedge  n  \leq 0.5 \wedge  o  \leq 1.0)$	30	15/8	3	[400, 300, 200]
Oscillator	$\square_{[0,300]} (a < 0.05)$	1	18/2	3	[280, 240, 200]

end-to-end systems that feature one single AI controller using raw input signals to generate control signals. We introduce the 8 control systems as follows:

- **CartPole:** This system consists of a pole attached to an unactuated joint connected to a cart moving along a track. Safety is ensured if the pole's angle  $\eta$  is less than  $90^\circ$  from the upright position, the cart distance  $\delta$  from the origin is within 0.3 meters, and the velocities  $v_1$  and  $v_2$  are below 0.3 and 0.5 m/s, respectively.
- **Pendulum:** This system is to maintain a pendulum attached to an unactuated joint in an upright position. Safety is ensured if the pendulum's angle  $\eta$  does not exceed the threshold of  $90^\circ$  or the pendulum's angular velocity  $\omega$  remains below  $\pi/2$  rad/s.
- **Quadcopter:** The control system aims to achieve stable flight for a quadcopter by limiting the angle  $\eta_1$  of the pitch rotation and the angle  $\eta_2$  of the yaw rotation not to exceed  $\pi/2$ .
- **Self Driving:** Discussed in detail in Section 3.
- **Lane Keeping:** This system aims to keep a vehicle between lane markers and centered in a possibly curved lane. Safety is ensured if a safe distance  $d$  to the cars in front is less than 0.9 meters.
- **4-Car Platoon and 8-Car Platoon:** These two systems model 4 or 8 vehicles forming a platoon and maintaining a safe relative distance between them. Each variable in the specification is the distance between two consecutive vehicles in the platoon, and safety is assured if they are below certain thresholds.
- **Oscillator:** This system consists of a two-dimensional switched oscillator and a 16-order filter. Safety is achieved if the output signal  $a$  of the filter remains below the safe threshold of 0.05.

Note that the inputs for these systems adhere to specific predefined ranges. The training data of AI controllers as well as the falsifying inputs are both from those ranges. The other falsifying inputs, which are outside the input ranges, are not considered within the scope of our work. Moreover, these control systems exhibit significant variation in the number of sub-specifications, ranging from 1 to 30 as shown in Table 2. This diversity allows for a comprehensive evaluation. Detailed information regarding the input ranges and sub-specifications for each control system can be found in our replication package, accessible through the link provided in Section 9.

We utilized the well-trained AI controllers for the 8 control systems from the official implementation of Zhu et al. [134], where these controllers were trained by Deep Deterministic Policy Gradient [60], a popular DRL algorithm for continuous control tasks. Table 2 shows the numbers of input and output variables of the AI controllers as well as their neural network sizes. The AI controllers have a wide range of sizes, featuring 2 or 3



layers with neuron counts varying from 440 (240+200) to 1200 (500+300+400); such sizes are sufficiently large for AI controllers (e.g., those controllers used in [100, 131] have no more than 400 neurons).

*5.1.2 Baseline.* We compare SYNTHIFY to PSY-TALiRO [98], a state-of-the-art and open-source tool for falsification of control systems. PSY-TALiRO recently won the 2022 ARCH-COMP prize for its technical achievements and excellent performance in falsification [30]. It is the Python version of a MATLAB tool called S-TaLiRo [6], which has been classified as the most mature tool ready for industrial deployment [52, 102]. In our study, we use this Python version of PSY-TALiRO to ensure consistency with the programming language of the Python-developed control systems we selected. PSY-TALiRO provides built-in optimization algorithms, including Uniform Random and Simulated Annealing. We chose the Simulated Annealing-based algorithm to support PSY-TALiRO in our experiments, as it has been shown to be more effective than Uniform Random [30, 93].

In our evaluation, we did not include several other open-source black-box falsification tools, such as Breach [22], ARISTEO [68], FalStar [31], ForeSee [129], zlscheck [45], NNV [100], and Mosaic [115]. These tools are specialized for MATLAB/Simulink models and written in different programming languages such as MATLAB, Java, Scala, or OCaml. The features of different programming languages can greatly impact the experimental results. Also, integrating these tools into our evaluation could pose significant compatibility challenges with the Python-developed control systems and toolchains used in our experiments. We believe that comparing SYNTHIFY to PSY-TALiRO, the only existing falsification toolbox natively in Python, is fair and appropriate since PSY-TALiRO is a state-of-the-art and well-established falsification tool that has been well studied. The 2022 ARCH-COMP prize awarded to PSY-TALiRO also demonstrates its advancement and maturity in falsification [30].

## 5.2 Implementation and Hyperparameters

SYNTHIFY is built on top of the implementation of Simulated Annealing (SA)-based falsification algorithm and other infrastructure components from PSY-TALiRO (version 1.0.0b3). To compute the quantitative robustness semantics, we integrate an off-the-shelf tool called RTAMT [71] (version 0.3.0) into our framework. We set the maximum number of iterations for the SA-based algorithm to 100, a recommended value by recent literature [68, 118]. By reusing all other default hyperparameters of PSY-TALiRO in the SA-based algorithm, we ensure a fair comparison while also evaluating the unique design of SYNTHIFY. For the Evolution Strategy (ES)-based synthesis algorithm, we use a population size of 50, and both the maximum number of iterations and time steps are set to 100. The noise standard deviation and the learning rate are set to 0.1 and 0.05, respectively. The threshold of the fitness is set to 0.1. The refinement follows the same hyperparameters as the ES-based synthesis algorithm, except that the learning rate is set to  $1e-3$  to avoid overfitting. In our  $\epsilon$ -greedy sampling strategy, we set the probability threshold  $\epsilon$  to 0.9. All the experiments are conducted using an Ubuntu 18.04 server equipped with an Intel Xeon E5-2698 CPU and 504 GB RAM.

## 5.3 Experimental Setup and Metrics

For a comprehensive evaluation, we employ two settings of the budgets of falsification in our experiments, as follows:

- **Setting 1:** Following the common practice of the annual ARCH workshop [28, 29], we specify the budget as the maximum number of falsification trials, which is set to 50.
- **Setting 2:** Following the common practice of search-based testing [51, 75, 135], we specify the budget as a fixed duration of time, which is set to 10 minutes, in line with prior studies [51, 92].

To mitigate the impact of randomness, we repeat each experiment 10 times and report the average results along with their statistical significance and effect size as suggested by Arcuri and Briand [7].

Then for each setting, we compare the performance of SYNTHIFY and PSY-TALiRO using the following metrics:

Table 3. Falsification results of SYNTHIFY and PSY-TALiRo on the 8 control systems. The  $\hat{A}_{12}$  column shows the effect size, where  $N$ ,  $M$ , and  $L$  indicate negligible, medium, and large effect sizes, respectively.

System	Setting 1 (50 trials)				Setting 2 (10 minutes)			
	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$
CartPole	99.8%	<b>100.0% (+0.2%)</b>	0.36	$N$ (0.55)	99.8	<b>561.1 (5.6×)</b>	$1.82e-4 < 0.01$	$L$ (1.00)
Pendulum	0.0%	<b>100.0% (−)</b>	$1.59e-5 < 0.01$	$L$ (1.00)	0.0	<b>23.9 (−)</b>	$4.84e-5 < 0.01$	$L$ (1.00)
Quadcopter	99.0%	<b>100.0% (+1.0%)</b>	0.07	$M$ (0.65)	82.4	<b>177.1 (2.1×)</b>	$5.20e-4 < 0.01$	$L$ (0.98)
Self-Driving	96.2%	<b>100.0% (+4.0%)</b>	$7.25e-4 < 0.01$	$L$ (0.90)	59.9	<b>150.9 (2.5×)</b>	$1.86e-4 < 0.01$	$L$ (0.80)
Lane Keeping	0.0%	<b>100.0% (−)</b>	$1.59e-5 < 0.01$	$L$ (1.00)	0.0	<b>159.3 (−)</b>	$6.20e-5 < 0.01$	$L$ (1.00)
4-Car Platoon	42.0%	<b>100.0% (+138.1%)</b>	$6.25e-5 < 0.01$	$L$ (1.00)	5.7	<b>626.6 (109.9×)</b>	$1.82e-4 < 0.01$	$L$ (1.00)
8-Car Platoon	0.0%	<b>100.0% (−)</b>	$1.59e-5 < 0.01$	$L$ (1.00)	0.0	<b>472.0 (−)</b>	$6.39e-5 < 0.01$	$L$ (1.00)
Oscillator	99.0%	<b>100.0% (+1.0%)</b>	$1.37e-2 < 0.05$	$L$ (0.75)	90.1	<b>470.7 (5.2×)</b>	$1.82e-4 < 0.01$	$L$ (1.00)
Average	54.5%	<b>100.0% (+83.5%)</b>	—	—	42.2	<b>330.2 (7.8×)</b>	—	—

- **Setting 1:** The effectiveness of each tool is evaluated by the falsification success rate, which is the ratio of the number of trials that uncover violations to 50 trials. This metric is commonly adopted in falsification studies [28, 30, 131]. The efficiency of each tool is measured by the falsification time, which is the time required to complete 50 falsification trials.
- **Setting 2:** The effectiveness of each tool is evaluated using the metric of the number of violations found by each tool within the given time budget of 10 minutes, following prior search-based testing studies [51, 95]. The efficiency of each tool is measured by the time required to reveal one safety violation.

Additionally, we also evaluate the sub-specification violation coverage of each tool, which is already defined in Section 3. Note that we report the performance improvements of SYNTHIFY over PSY-TALiRo in terms of the relative percentage difference in metrics. This is calculated by taking the difference between the two tools' performance and dividing it by the performance of PSY-TALiRo.

## 6 EVALUATION RESULTS

### 6.1 RQ1: Effectiveness of SYNTHIFY

Table 3 shows the falsification results of SYNTHIFY and PSY-TALiRo for the 8 control systems. As mentioned in Section 5, we repeat each experiment 10 times and report the average results here. For Setting 1, we compare the falsification success rates of the two tools and find that SYNTHIFY outperforms PSY-TALiRo on all control systems. SYNTHIFY achieves a 100% falsification success rate on all control systems. In contrast, PSY-TALiRo only achieves an average success rate of 54.5%. On average, SYNTHIFY outperforms PSY-TALiRo by 83.5% ( $(100\% - 54.5\%)/54.5\% = 83.5\%$ ) in terms of the falsification success rate.

To account for the impact of randomness and ensure the credibility of the results, we used the Mann-Whitney U test [65], which is a non-parametric statistical test used to compare two independent groups. We use it to measure the statistical significance of the difference in falsification success rates between the two tools across the 10 experiments. Our results show that SYNTHIFY significantly outperforms PSY-TALiRo on 6 out of 8 control systems, with a confidence level of 99% or 95%. On CartPole and Quadcopter, the improvements achieved by SYNTHIFY are not as significant because the falsification success rates of PSY-TALiRo are already quite high.

We also use the Vargha-Delaney statistic  $\hat{A}_{12}$  [104] to measure the effect size. This non-parametric measure gives the probability that a randomly chosen value from one group is higher or lower than one from another group. It is commonly used to evaluate search-based testing methods due to their randomness [67, 73, 74]. According to Vargha and Delaney [104],  $\hat{A}_{12}$  values greater than 0.71 (or less than 0.29) indicate a “large” effect size. Values between 0.64 and 0.71 (or between 0.29 and 0.36) correspond to a “medium” effect size, while values between 0.36 and 0.64 suggest a “negligible” effect size. As shown in Table 3, SYNTHIFY significantly outperforms PSY-TALiRo

Table 4. Falsification time of SYNTHIFY and PSY-TALiRo. The  $\hat{A}_{12}$  column shows the effect size,  $L$  indicates a large effect size.

System	Total falsification time over 50 trials				Time required to reveal one violation			
	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$
CartPole	286.5s	<b>88.2s (3.2×)</b>	1.82e-4<0.01	$L$ (1.00)	6.1s	<b>1.1s (5.5×)</b>	2.83e-3<0.01	$L$ (1.00)
Quadcopter	127.9s	<b>46.6s (2.7×)</b>	1.83e-4<0.01	$L$ (1.00)	8.3s	<b>4.1s (2.0×)</b>	2.80e-4<0.01	$L$ (1.00)
Self-Driving	543.4s	<b>182.8s (3.0×)</b>	1.83e-4<0.01	$L$ (1.00)	15.4s	<b>5.5s (2.8×)</b>	8.62e-4<0.01	$L$ (0.97)
4-Car Platoon	1462.7s	<b>80.8s (18.1×)</b>	1.83e-4<0.01	$L$ (1.00)	129.2s	<b>1.0s (129.2×)</b>	3.30e-4<0.01	$L$ (0.98)
Oscillator	295.1s	<b>86.5s (3.4×)</b>	1.83e-4<0.01	$L$ (1.00)	6.8s	<b>1.3s (5.2×)</b>	1.31e-3<0.01	$L$ (0.93)
Average	543.1s	<b>97.0s (5.6×)</b>	—	—	33.2s	<b>2.6s (12.8×)</b>	—	—

with a large effect size on 6 out of 8 control systems. Only on CartPole and Quadcopter, we observe negligible and medium effect sizes. But overall, the results show that SYNTHIFY is more effective than PSY-TALiRo in finding violations of safety specifications.

In Setting 2, SYNTHIFY also consistently outperforms PSY-TALiRo across all control systems. Table 3 presents the results of Setting 2, where SYNTHIFY and PSY-TALiRo are compared in terms of the number of violations they reveal within the same time budget of 10 minutes. The results show that, on average, SYNTHIFY can uncover 330.2 violations of safety specifications within 10 minutes, which is 7.8 times more than PSY-TALiRo (42.3 violations). This result highlights SYNTHIFY’s ability to reveal significantly more safety violations than PSY-TALiRo within the same time budget. Furthermore, the results of Mann-Whitney U tests and  $\hat{A}_{12}$  statistics show that SYNTHIFY significantly outperforms PSY-TALiRo with a 99% confidence level and a large effect size.

**Answers to RQ1:** SYNTHIFY significantly outperforms PSY-TALiRo in terms of effectiveness. Within the same falsification trials, SYNTHIFY outperforms PSY-TALiRo by 83.5% in terms of falsification success rate. Within the same time budget, SYNTHIFY reveals 7.8× more safety violations than PSY-TALiRo.

## 6.2 RQ2: Efficiency of SYNTHIFY

The comparison results of the falsification time over the 50 trials of Setting 1 are shown in Table 4. Note that the falsification time of SYNTHIFY also includes the time for the synthesis and refinement process. Pendulum, Lane Keeping, and 8-Car Platoon are excluded from this comparison as PSY-TALiRo failed to falsify their safety specifications in our RQ1 experiments. For the remaining five control systems, we observe that SYNTHIFY consistently outperforms PSY-TALiRo in terms of falsification time for all of them. On average, SYNTHIFY is 5.6× faster than PSY-TALiRo in completing 50 falsification trials. These performance differences are statistically significant at the 99% confidence level, as confirmed by the Mann-Whitney U tests. Additionally, the  $\hat{A}_{12}$  statistics indicate that the effect sizes of the improvements are large in all cases. Thus, SYNTHIFY offers significant efficiency gains in the falsification process across different control systems.

In Setting 2, SYNTHIFY also takes much less time to find a single violation compared to PSY-TALiRo. Table 4 shows the time required for each tool to find a violation, still excluding a few control systems due to that PSY-TALiRo cannot perform any falsification on them. On average, SYNTHIFY takes only 2.6 seconds to find one single violation, which is 12.8× faster than PSY-TALiRo in finding one single violation. Still, the statistical differences between the time taken by the two tools indicate that SYNTHIFY significantly outperforms PSY-TALiRo across different control systems at the 99% confidence level. The  $\hat{A}_{12}$  statistics also show that SYNTHIFY significantly outperforms PSY-TALiRo with a large effect size.

Table 5. Sub-specification violation coverage of SYNTHIFY and PSY-TALiRo. The  $\hat{A}_{12}$  column shows the effect size where  $N$  and  $L$  indicate negligible and large effect size, respectively.

Control System	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$
CartPole	50.0%	<b>98.6% (+97.2%)</b>	$2.42e-5 < 0.01$	$L (1.00)$
Pendulum	0.0%	<b>25.0% (-)</b>	$1.59e-5 < 0.01$	$L (1.00)$
Quadcopter	75.0%	<b>100.0% (+33.3%)</b>	$1.59e-5 < 0.01$	$L (1.00)$
Self-Driving	50.0%	<b>100.0% (+100.0%)</b>	$9.67e-3 < 0.01$	$L (0.80)$
Lane Keeping	0.0%	<b>90.0% (-)</b>	$3.29e-5 < 0.01$	$L (1.00)$
4-Car Platoon	7.1%	<b>76.4% (+976.1%)</b>	$4.14e-5 < 0.01$	$L (1.00)$
8-Car Platoon	0.0%	<b>81.0% (-)</b>	$4.17e-5 < 0.01$	$L (1.00)$
Oscillator	100%	100%	1.00	$N (0.50)$
Average	35.3%	<b>83.9% (+137.7%)</b>	—	—

**Answers to RQ2:** SYNTHIFY significantly outperforms PSY-TALiRo in terms of falsification time and the time required to find a violation. Specifically, SYNTHIFY is 5.6× faster than the baseline at completing 50 falsification trials. Moreover, SYNTHIFY can find a violation 12.8× faster than PSY-TALiRo.

### 6.3 RQ3: Sub-specification Violation Coverage

SYNTHIFY achieves an average sub-specification violation coverage of 83.9%, which is a substantial improvement of 137.7% over PSY-TALiRo’s coverage of 35.3%, as listed in Table 5. Note that the coverage results are consistent for both Setting 1 and Setting 2. Of all the control systems, only the results on Oscillator show no improvement. The reason for this is that Oscillator contains only one sub-specification, and both SYNTHIFY and PSY-TALiRo can effectively falsify it. However, for all other control systems, SYNTHIFY outperforms PSY-TALiRo by at least 33.3% in terms of sub-specification violation coverage.

Similar to RQ1 and RQ2, we perform Mann-Whitney U tests and  $\hat{A}_{12}$  statistics, and the results confirm that SYNTHIFY significantly outperforms PSY-TALiRo on all control systems except Oscillator, with a 99% confidence level and a large effect size. These results further highlight the superiority of SYNTHIFY in achieving higher sub-specification violation coverage than PSY-TALiRo.

**Answers to RQ3:** SYNTHIFY significantly outperforms PSY-TALiRo in terms of sub-specification violation coverage. Specifically, SYNTHIFY can falsify 83.9% of safety specifications on average, which is a 137.7% improvement over PSY-TALiRo.

## 7 ADDITIONAL ANALYSIS

Aside from the main results presented in Section 6, we provide additional analyses to further understand the performance of SYNTHIFY. These analyses include an ablation study to assess the importance of each core module in SYNTHIFY, a breakdown of the running time for each process in SYNTHIFY, an evaluation of the alignment between the proxy programs and AI controllers, and an investigation of SYNTHIFY’s performance on control systems with large AI controllers. We also detail the selection for the hyperparameters of our sub-specification sampling strategy, offer a discussion on the practical implications of our work in the era of end-to-end AI-enabled control systems, as well as the limitations of our study.

Table 6. Ablation study of core modules in SYNTHIFY. Results are averaged over 50 falsification trials on 8 control systems.

	Success Rate	Time	Coverage
SYNTHIFY	100.0%	97.0s	83.9%
w/o Refinement	87.5%	94.5s	71.4%
w/o Synthesis & Refinement	51.7%	781.3s	53.1%
w/o $\epsilon$ -greedy Sampling	100.0%	78.1s	46.0%

Table 7. Average running time of each process in SYNTHIFY on 8 control systems for 50 falsification trials.

Synthesis	Refinement	Proxy Program
9.1s / 9.4%	1.4s / 1.4%	4.2s / 4.3%
AI Controller	Plant Execution	Algorithm
5.0s / 5.2%	64.3s / 66.3%	13.0s / 13.4%

## 7.1 Ablation Study

We perform an ablation study by comparing the performance of SYNTHIFY with different variants. Table 6 shows the average results over 50 falsification trials on the 8 control systems. The variants considered in the study are as follows:

**w/o Refinement.** SYNTHIFY without the refinement process (described in Section 4.3). Without this step, the synthesized proxy program may not behave sufficiently similarly to the AI controller, impacting the effectiveness of SYNTHIFY. The success rate of SYNTHIFY in finding safety violations within 50 falsification trials drops from 100.0% to 87.5%, and the sub-specification violation coverage is also reduced by 12.5% (from 83.9% to 71.4%). These results confirm that the refinement process improves the effectiveness of SYNTHIFY.

**w/o Synthesis & Refinement.** SYNTHIFY without the synthesis (described in Section 4.1) and refinement (Section 4.3) processes. These two steps are designed to improve the scalability of falsification. Without them, the falsification time increases sharply by  $8.1\times$  (from 97.0s to 781.3s), and both the success rate and sub-specification violation coverage greatly decrease. These results highlight the importance of the synthesis and refinement processes in achieving efficient and effective falsification.

**w/o  $\epsilon$ -greedy Sampling.** SYNTHIFY without the  $\epsilon$ -greedy sampling strategy (described in Section 4.2). We directly use the conjunctive specification in falsification but not sample from it. SYNTHIFY achieves an average sub-specification violation coverage of only 46.0%, which is 37.9% (from 83.9% to 46.0%) lower. This result shows the importance of the  $\epsilon$ -greedy sampling in improving sub-specification coverage.

These results highlight the importance of each core process of SYNTHIFY in achieving efficient and effective falsification.

## 7.2 Analysis of Running Time

Table 7 shows the average running time of each process in SYNTHIFY on 8 control systems for 50 falsification trials. The synthesis process takes only 9.1 seconds on average, accounting for 9.4% of the total running time. The refinement process takes even less time, averaging 1.4 seconds, or 1.4% of the total running time. These results show that the synthesis and refinement are both lightweight and do not significantly impact the overall running time of SYNTHIFY. The execution of the proxy program and the AI controller take 4.3% and 5.2% of the total running time, respectively, neither of which is significant. This result clearly shows AI controller is no longer the bottleneck of the running time in falsification as described in Section 3, and the synthesis of a computationally



cheaper proxy program has led to significant improvements in efficiency, shifting the bottleneck to the execution of the plant, which can be future work to further improve the efficiency of SYNTHIFY.

### 7.3 Alignment between Proxy Programs and AI Controllers

SYNTHIFY’s effectiveness relies on the functional alignment between the synthesized proxy programs and AI controllers. Intuitively, the closer the proxy programs match the behavior of the AI controllers, the more likely they are to produce similar outputs, with fewer spurious violations during the refinement process—i.e., if the AI controllers produce unsafe outputs that result in safety violations, then the proxy programs are likely to do the same. This section conducts an analysis to evaluate this alignment. For this evaluation, we synthesize proxy programs for the 8 control systems using the same experiment settings as described in Section 5. Both the proxy programs and AI controllers run for 1000 time steps to gather their outputs. Additionally, we run the falsification for 50 trials and track the number of spurious violations encountered during refinement.

We first use the mean absolute error (MAE) to directly quantify the difference between the outputs of the AI controllers and the proxy programs. A lower MAE suggests stronger alignment. Since the AI controllers produce continuous outputs rather than discrete labels, the accuracy metric used in [109] for label-based predictions is not suitable for this comparison. MAE, which is commonly used to evaluate regression models [12], is more appropriate in this context. To ensure comparability, we normalize the outputs of both the AI controllers and the proxy programs to the range [0, 1], as their scales differ across various control systems. After normalization, we calculate the MAE using the following formula:

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_{\text{AI}}^{(i)} - y_{\text{proxy}}^{(i)}| \quad (7)$$

where  $N$  represents the number of time steps, and  $y_{\text{AI}}^{(i)}$  and  $y_{\text{proxy}}^{(i)}$  denote the normalized outputs of the AI controller and proxy program at time step  $i$ , respectively. We repeat the experiments 10 times and present the average results in Table 8. The MAE for the outputs averages 0.19 across all control systems, indicating that the proxy programs’ outputs differ by only 19% (as the outputs are already normalized to the [0, 1] range) from those of the AI controllers. This difference is considered small and acceptable by the prior studies [38, 39] that also evaluate the alignment between the outputs of real systems and their proxies. These results demonstrate a strong functional alignment between the proxy programs and the AI controllers.

Beyond comparing outputs, we also evaluate alignment between the AI controllers and proxy programs by examining the number of spurious violations during the refinement process, as shown in Table 8. Our results show that, on average, SYNTHIFY identifies 2.1 spurious violations for the purpose of refining the proxy programs. The minimal occurrence of spurious violations further demonstrates the proxy programs’ effectiveness in capturing the AI controllers’ behavior. This strong alignment provides a solid foundation for applying falsification techniques on the proxy programs to identify safety violations in AI-enabled control systems.

### 7.4 Handling Large AI Controllers

As mentioned in Section 5.1, the AI controllers used in our evaluations vary significantly in size, with neuron counts ranging from 440 to 1200, which are relatively large compared to those used in other studies like [100, 131], providing a diverse basis for evaluation. However, we anticipate that larger models will be integrated into control systems as hardware continues to advance. Therefore, we conduct additional experiments to evaluate SYNTHIFY’s performance on control systems with larger AI controllers. These new experiments use three variant models trained in the CartPole, Pendulum, and Self-Driving control systems, also released by Zhu et al. [134], which feature larger AI controllers with 3 layers containing 1200, 900, and 800 neurons, respectively—an increase of 2.4× to 6.6× the number of neurons compared to the 8 controllers used in our main experiments.



Table 8. The mean absolute error (MAE) of the outputs of the AI controllers and the synthesized proxy programs, and the number of spurious violations used in the refinement. Results are averaged over experiments repeated 10 times.

Control System	MAE (Outputs)	# Spurious Violations
CartPole	0.31	1.6
Pendulum	0.14	4.4
Quadcopter	0.10	6.3
Self-Driving	0.07	3.5
Lane Keeping	0.22	1.0
4-Car Platoon	0.11	0
8-Car Platoon	0.18	0
Oscillator	0.40	0
Average	0.19	2.1

Table 9. Falsification results of SYNTHIFY and PSY-TALiRo on three control systems with large AI controllers. Relative improvements are shown in parentheses. The  $\hat{A}_{12}$  column shows the effect size, where  $N$ ,  $M$ , and  $L$  indicate negligible, medium, and large effect sizes, respectively.

System	Setting 1 (50 trials)				Setting 2 (10 minutes)			
	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$
CartPole (Large)	17.6%	<b>100.0% (+468.2%)</b>	7.49e-3<0.01	$L$ (1.00)	15.6	<b>818.5 (52.5×)</b>	1.83e-3<0.01	$L$ (1.00)
Pendulum (Large)	0.0%	<b>100.0% (–)</b>	3.91e-3<0.01	$L$ (1.00)	0.0	<b>49.3 (–)</b>	1.82e-5<0.01	$L$ (1.00)
Self-Driving (Large)	100.0%	100.0%	1.0	$N$ (0.50)	65.0	<b>298.9 (4.6×)</b>	9.34e-5<0.01	$L$ (1.00)
Average	39.2%	<b>100.0% (+155.1%)</b>	–	–	26.9	<b>388.9 (14.5×)</b>	–	–

Table 10. Sub-specification violation coverage of SYNTHIFY and PSY-TALiRo. The  $\hat{A}_{12}$  column shows the effect size where  $N$  and  $L$  indicate negligible and large effect size, respectively.

Control System	PSY-TALiRo	SYNTHIFY	$p$ -value	$\hat{A}_{12}$
CartPole (Large)	12.5%	<b>100.0% (+700.0%)</b>	3.9e-3<0.01	$L$ (1.00)
Pendulum (Large)	0.0%	<b>100.0% (–)</b>	3.7e-4<0.01	$L$ (1.00)
Self-Driving (Large)	50.0%	<b>100.0% (+100.0%)</b>	3.9e-3<0.01	$L$ (0.80)
Average	31.3%	<b>100.0% (+219.5%)</b>	–	–

All experiments are conducted under the same settings described in Section 5, and the results are presented in Table 9. SYNTHIFY consistently outperforms the baseline in both settings, achieving a relative improvement of 155.1% in the falsification success rate across 50 trials and 14.5× more violations found within 10 minutes. Statistical tests confirm the significance and effect size of these improvements, with  $p$ -values below 0.01 and large effect sizes (except for the Self-Driving system, where both our SYNTHIFY and PSY-TALiRo achieve a 100% success rate). Additionally, sub-specification violation coverage is significantly improved compared to PSY-TALiRo, with an average relative increase of 219.5%, as shown in Table 10. These results demonstrate that SYNTHIFY remains highly effective at falsifying safety violations in control systems with large AI controllers, underscoring its potential for future applications in more complex systems.

Table 11. Sub-specification violation coverage of SYNTHIFY with different hyperparameters of the  $\epsilon$ -greedy sampling strategy.

System	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
CartPole	50.0%	75.0%	50.0%	50.0%	100%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Pendulum	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%	25.0%
Quadcopter	75.0%	75.0%	75.0%	75.0%	50.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Self-Driving	25.0%	25.0%	25.0%	25.0%	75.0%	75.0%	100.0%	75.0%	100.0%	100.0%	100.0%
Lane Keeping	50.0%	50.0%	50.0%	50.0%	50.0%	100.0%	100.0%	50.0%	100.0%	100.0%	50.0%
4-Car Platoon	71.4%	85.7%	92.9%	92.9%	92.9%	100.0%	92.9%	92.9%	100.0%	100.0%	92.9%
8-Car Platoon	73.3%	86.7%	83.3%	80.0%	76.7%	86.7%	83.3%	86.7%	83.3%	90.0%	76.7%
Oscillator	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Average	59.8%	68.3%	68.3%	67.3%	73.3%	86.7%	90.0%	80.0%	90.0%	93.3%	83.3%

Table 12. Falsification results on Self Driving control system with two cases of intricate safety specifications.

Safety Specification	Setting 1 (50 trials)		Setting 2 (10 minutes)		Coverage	
	PSY-TALiRo	SYNTHIFY	PSY-TALiRo	SYNTHIFY	PSY-TALiRo	SYNTHIFY
$\square_{[0,200]} \diamond_{[0,20]} ( \eta  < 90^\circ \wedge  d  < 2.0)$	100.0%	100.0%	98.8	<b>378.8 (3.8×)</b>	50.0%	<b>100.0% (+100.0%)</b>
$\square_{[0,200]} ( d  > 2.0 \rightarrow \diamond_{[0,20]} ( d  < 2.0))$	66.7%	<b>100.0% (+49.9%)</b>	38.4	<b>154.5 (4.0×)</b>	50.0%	<b>100.0% (+100.0%)</b>
Average	83.4%	<b>100.0% (+19.9%)</b>	68.6	<b>266.7 (3.9×)</b>	50.0%	<b>100.0% (+100.0%)</b>

## 7.5 Hyperparameters of the Sampling Strategy

Our  $\epsilon$ -greedy sampling strategy (Section 4.2) is governed by the hyperparameter  $\epsilon$ , which controls the balance between selecting a sub-specification randomly (exploration) and selecting the most promising one (exploitation). To find the optimal value of  $\epsilon$ , we use grid search, a common method for hyperparameter tuning [10]. Grid search evaluates performance for each hyperparameter within a predefined set of options. In our experiments, we vary  $\epsilon$  within the range [0.0, 1.0] in increments of 0.1, while keeping all other hyperparameters constant (i.e., following the settings described in Section 5). A value of 0.0 corresponds to a fully greedy strategy, always selecting the most promising sub-specifications, while 1.0 corresponds to a fully exploratory strategy, selecting sub-specifications randomly. We conduct 50 falsification trials and evaluate SYNTHIFY’s sub-specification violation coverage across 8 control systems. The results, shown in Table 11, indicate that SYNTHIFY’s performance improves with increasing  $\epsilon$  values, with an optimal value of 0.9 yielding the highest average sub-specification violation coverage of 93.3%. Consequently, we set  $\epsilon$  to 0.9 in our experiments by default. However, the optimal value may vary depending on the control system and safety specification, necessitating further tuning in practice.

## 7.6 Falsifying Intricate Safety Specifications

In our experiments, the conjunctive safety specifications typically take the form of  $\square_{\mathcal{I}} \varphi_1 \wedge \varphi_2$ , requiring the system to satisfy both  $\varphi_1$  and  $\varphi_2$  at every time step within the interval  $\mathcal{I}$ . However, real-world safety specifications are often more intricate, involving nested temporal operators or complex logical expressions. We believe SYNTHIFY is capable of falsifying such complex safety specifications, benefiting from our implementation on top of the infrastructure components from PSY-TALiRo and utilizing the off-the-shelf tool RTAMT to compute the qualitative robustness semantics (as we mentioned in Section 5.2). Both can handle complex safety specifications, equipping SYNTHIFY with the same capability.

To further quantitatively assess SYNTHIFY’s effectiveness in falsifying more intricate safety specifications, we construct additional specifications that involve nested temporal operators, such as  $\square_{\mathcal{I}} (\diamond_{\mathcal{I}} \varphi)$  or  $\square_{\mathcal{I}} (\varphi_1 \rightarrow \diamond_{\mathcal{I}} \varphi_2)$ , for more experiments. Since the benchmarks used in our experiments only include specifications of the form

$\Box_I \varphi$ , these additional specifications can be only manually designed, requiring domain expertise and making them difficult to scale. We only evaluate SYNTHIFY on two intricate safety specifications using the Self Driving control system introduced in Section 3, which is representative of real-world autonomous driving systems, as a proof of concept.

The safety specifications and corresponding results are presented in Table 12. The first specification ( $\Box_{[0,200]} \Diamond_{[0,20]} (|\eta| < 90^\circ \wedge |d| < 2.0)$ ) requires the car to reach a state where the steering angle  $\eta$  is within 90 degrees and the distance  $d$  between the car’s front and the road’s centerline is within 2.0 meters, within 20 time steps at every point during the 200 time steps, enforcing a proactive condition where the system must regularly check and ensure safety. The second specification ( $\Box_{[0,200]} (|d| > 2.0 \rightarrow \Diamond_{[0,20]} (|d| < 2.0))$ ) requires the car to steer back to a safe distance within 20 time steps after  $d$  exceeds 2.0 meters during the 200 time steps, enforcing a reactive condition where the system must respond to unsafe situations. We evaluate SYNTHIFY on these specifications using the same experimental settings as in Section 5 and the same evaluation metrics. The results show that SYNTHIFY is highly effective in falsifying these complex safety specifications, achieving a 100% success rate across all case studies, which is 19.9% higher than the baseline. Additionally, the number of uncovered violations increased by 3.9× within 10 minutes compared to the baseline. The sub-specification violation coverage also improved significantly, with an average relative improvement of 100%. These results highlight SYNTHIFY’s ability to handle more complex safety specifications and its effectiveness beyond the simpler conjunctive specifications used in the primary experiments, demonstrating its potential for real-world applications.

## 7.7 Potential Applications in End-to-End AI Systems

The adoption of end-to-end AI-enabled control systems has recently surged in various domains, such as autonomous driving [16] and robotics [35]. These systems use a single neural network to map raw sensor data directly to control signals, achieving enhanced performance, as exemplified by real-world systems like Tesla FSD [97] and NVIDIA DRIVE [72]. We suggest that SYNTHIFY has been developed in line with the trend of adopting end-to-end AI systems, and quantitatively evaluated through experiments on 8 end-to-end AI-enabled control systems. A further qualitative discussion supporting SYNTHIFY’s effectiveness in this context is provided below.

First and foremost, SYNTHIFY enhances the efficiency and sub-specification coverage of falsification for end-to-end AI systems. As mentioned in Section 1, these systems typically involve complex AI controllers with numerous parameters for processing inputs and generating control signals in one go [11, 49]. This complexity leads to high computational costs during testing [36, 75]. SYNTHIFY mitigates this by synthesizing efficient proxy programs that replace the AI controllers, significantly reducing the runtime of falsification. Moreover, thorough testing of end-to-end AI systems requires diverse test inputs [40, 61]. SYNTHIFY meets this need by maximizing sub-specification coverage of the conjunctive safety specification, as a more diverse set of test inputs is likely to falsify more sub-specifications [106, 126]; intuitively vice versa. The  $\epsilon$ -greedy sampling strategy employed by SYNTHIFY promotes the exploration of various sub-specifications, enabling the falsification process to uncover diverse safety violations by targeting different aspects of the safety requirements.

Furthermore, certain modules in SYNTHIFY, such as our synthesis and  $\epsilon$ -greedy sampling strategy, can complement existing verification and testing tools for end-to-end AI systems. For instance, our synthesis module could be applied to generate simple programs that closely approximate the behavior of complex AI controllers but are more amenable to formal verification methods [32, 58], like constructing barrier certificates [134] or performing reachability analysis [99]. In the context of testing, our  $\epsilon$ -greedy sampling strategy can be integrated into existing falsification tools like [24, 131] to strengthen their ability to falsify conjunctive specifications. Alternatively, it can support automated testing frameworks that handle multiple objectives [3, 36], helping developers efficiently prioritize and focus on the most promising objectives during testing.

Beyond testing and verification, SYNTHIFY aids developers in inspecting end-to-end AI systems for interpretability and explainability—two essential factors in ensuring the trustworthiness of AI-enabled systems [33]. Interpretability focuses on understanding the system’s inherent characteristics. Developers can leverage SYNTHIFY’s identified safety violations and corresponding falsifying inputs to derive formal specifications and properties of AI controllers [34, 81], serving as proof of the system’s safety (or lack thereof). Explainability, on the other hand, involves using external proxies to clarify the system’s operations [107]. SYNTHIFY’s synthesized proxy programs, with straightforward control logic, can provide explanations for generated control signals by transparently revealing the execution process. This allows developers to better understand the AI controller’s decision-making and identify potential vulnerabilities or biases within the system.

## 7.8 Limitations and Potential Solutions

While SYNTHIFY demonstrates success in several cases within our evaluation, its primary limitation lies in generalizing to control systems with more complex inputs, such as images or videos. These inputs often have many dimensions and correspond to numerous variables in our linear program sketch, which can significantly expand the search space during the synthesis process, making it difficult to identify a suitable proxy program that accurately captures the AI controller’s behavior. One potential solution is to use dimensionality reduction techniques like t-distributed Stochastic Neighbor Embedding (t-SNE) [42] to project the inputs into a lower-dimensional space before synthesizing the proxy program. Alternatively, we could move away from the current linear program sketch and instead synthesize a compact neural network capable of processing high-dimensional inputs using techniques like model compression [87–89]. Despite this limitation, our methods are currently well-suited for control systems with numerical inputs like scalar values or vectors, thus we leave the extension to more complex inputs as future work.

## 7.9 Threats to Validity

One threat to internal validity is the randomness of the falsification algorithms used in our experiments. To mitigate it, we repeat each experiment 10 times and use appropriate statistical tests to account for both statistical significance and effect size, as recommended by Arcuri and Briand [7].

In terms of external validity, one threat is that the results of our analysis may not be generalizable. To mitigate it, we have carefully selected a representative set of control systems with diverse tasks to evaluate our method and baseline under different settings, which ensures that our results are not biased. However, as these control systems take the input as numerical values, the generalization of our results to control systems with more complex inputs, such as images or videos, may need further investigation.

One potential threat to construct validity is that the evaluation metrics may not fully capture the performance of falsification tools. To mitigate it, we used a total of five commonly-used evaluation metrics that are introduced in Section 5.3 to compare the effectiveness and efficiency of SYNTHIFY and the baseline from a comprehensive set of perspectives.

## 8 RELATED WORK

This section presents an overview of studies relevant to this paper, including (1) falsification of control systems, (2) program synthesis for AI models, and (3) surrogate-assisted testing and verification, and (4) multi-objective optimization in AI-enabled control systems.

### 8.1 Falsification of AI-Enabled Control Systems

Falsification [4, 6, 15, 27, 31, 68, 102, 108] has emerged as an alternative to formal verification that provides better efficiency to uncover violations of a given safety specification. There comes many falsification techniques

supported by optimization algorithms, such as Hill Climbing [66], Simulated Annealing [1, 98], Monte Carlo Tree Search [128, 129], gradient-based methods [2, 116], as well as some machine learning-guided methods [5, 125].

To date, there has been limited attention directed toward the falsification of AI-enabled control systems, with only a few early attempts [25, 93, 115, 117, 131]. Among these studies, Song et al. [93] evaluate the reliability of AI-enabled control systems using two existing falsification tools [6, 22] and highlight that they need further improvements, which necessitates our work. Zhang et al. [131] and Yaghoubi et al. [117] develop gray-box falsification techniques that make different assumptions compared to our study. They assume the AI controllers or plants accessible and collect not only the input and output signals from the control system but also intermediate information of the AI controllers or plants as guidance for falsification. In comparison, our approach is a black-box falsification tool that solely relies on the input and output signals, making it more practical particularly in real-world applications where the controller or plants can be complex or even inaccessible if they are confidential products from third-party vendors.

VerifAI [25], NNV [100], and Mosaic [115] are all safety analysis frameworks for AI-enabled control systems, each incorporating a falsification module. However, there are key distinctions between these tools and our SYNTHIFY. VerifAI is designed for systems with AI models in the sensor components, while our work primarily concerns systems with AI models in the controller parts. In addition, both NNV and Mosaic require specific input specifications that differ from those supported by our SYNTHIFY. NNV relies on assert statements as safety specifications, while Mosaic necessitates probabilistic computation tree logic (PCTL) specifications. These particular specifications are not compatible with our benchmark systems, and our SYNTHIFY, like many other existing falsification tools, operates based on Signal Temporal Logic (STL) specifications [6, 22, 98]. Attempting to translate STL specifications into assert statements or PCTL specifications would introduce additional manual efforts, which are not only time-consuming but also prone to errors. Such translations also fall outside the scope of this paper. Consequently, in our comparison experiments, we have chosen to use PSY-TALiRo [98] as our baseline tool.

## 8.2 Surrogate-Assisted Falsification

Several studies have also focused on the safety analysis of control systems, though not specifically AI-enabled ones, and aligned with our work by leveraging surrogate models. Some of these studies approximate control systems with surrogate models trained on data collected from the systems, which can replace the original systems during falsification. For example, Menghi et al. [68] propose ARISTEO, which replaces computationally-intensive control systems with surrogate models derived from system identification techniques to improve falsification efficiency. Waga et al. [108] introduce FalCAuN, which learns a Mealy machine as a surrogate model and falsifies it using model checking, while Bak et al. [9] learn a Koopman model and falsify it through reachability analysis. Atanu et al. [56] propose to build a DNN model that approximates the behavior of the control system and use adversarial attack algorithms for falsification.

Another set of studies integrates surrogate models into the falsification process. For instance, several studies [20, 69, 77, 80, 127] employ Bayesian optimization with Gaussian processes as surrogate models to estimate the robustness semantics of STL specifications, guiding the search for safety violations. Additionally, Peltomäki et al. [76] propose training a surrogate DNN model to directly generate falsifying inputs for control systems.

While our synthesized proxy programs may resemble surrogate models, they fundamentally differ from those in above studies. First, in terms of form, our proxy programs implement linear controllers, whereas surrogate models in prior work are typically DNNs (e.g., [56, 76]) or machine learning models (e.g., [20, 77]). The simplicity of our programs offers greater efficiency. Second, regarding purpose, our proxy programs replicate the functionality of AI controllers, acting as their stand-ins to facilitate efficient falsification by minimizing the need to run the controllers themselves. This addresses the unique computational challenges associated with AI controllers, which



are expensive to execute due to their vast number of parameters. In contrast, surrogate models in earlier work are usually designed to replace control systems or predict robustness semantics. Third, in terms of underlying technique, our approach utilizes Evolution Strategy for program synthesis, while prior methods often rely on Bayesian optimization (e.g., [77, 80]) or DNNs (e.g., [56, 68, 76]) to build surrogate models.

These distinctions not only make our approach unique but also allow it to complement existing methods, further enhancing the efficiency of falsification in AI-enabled control systems. While combining our approach with existing surrogate-assisted falsification methods is a promising direction, a detailed evaluation is beyond the scope of this work and is left for future research.

### 8.3 Program Synthesis for AI models

Although program synthesis has been studied for decades [90], very little attention has been paid to program synthesis specialized for AI models. A few studies [101, 105] have proposed methods for generating programmatic policies from well-trained DRL models, which are more interpretable than neural networks. Additionally, Zhu et al. [134] introduce an inductive synthesis-based method that can synthesize deterministic and verified programs from DRL-based models. Inala et al. [44] and Qiu et al. [78] synthesize programmatic controllers to perform continuous control tasks. Our approach differs from these studies in that we propose an Evolution Strategy-based algorithm that can effectively synthesize programs as proxies for AI controllers, which we believe is both technically novel and offers a unique perspective on improving the efficiency of falsification in AI-enabled control systems.

### 8.4 Multi-Objective Optimization in AI-Enabled Control Systems

Our focus on covering various sub-specifications can be viewed as a multi-objective optimization problem. Thus, a relevant body of work lies in the multi-objective optimization within AI-enabled control systems, and we particularly mention those in the areas of testing, repair, and enhancement.

In the context of testing, Zolfagharian et al. [135] and Haq et al. [36, 37] propose multi-objective search-based testing methods for AI-enabled control systems, accounting for diverse fitness functions or safety requirements. Further studies address repairing these systems. Lyu et al. [62] repair AI controllers' unsafe behaviors by adding patches to control signals via multi-objective optimization, while MORTAR [110] replaces unsafe signals from AI controllers with repaired alternatives through gradient-based optimization. Both methods solve the multi-objective problem of maximizing satisfaction of safety requirements while minimizing changes to the original control signals. Regarding safety enhancement without directly repairing the violations, SIEGE [94] stands out by presenting an ensemble framework that synergistically combines multiple AI controllers to improve the system's reliability across multiple control objectives. GenSafe [133] corrects the behaviors of AI controllers that potentially deviate from safety constraints during training to enhance their safety while also maintaining performance at a satisfactory level. In addition to all of the above, Zhou et al. [132] establish a benchmark for multiple robotic manipulation tasks, providing a platform to evaluate the safety and effectiveness of AI controllers in a multi-objective context.

## 9 CONCLUSION AND FUTURE WORK

In this study, we propose SYNTHIFY, a novel falsification framework specifically designed for AI-enabled control systems. The key features of SYNTHIFY include the use of a synthesized program to serve as a proxy of the AI controller, and the application of the  $\epsilon$ -greedy strategy to sample a promising safety specification from the conjunctive requirement. We evaluated SYNTHIFY on 8 publicly-available control systems and compared it with a state-of-the-art falsification tool PSY-TALiRo. Our evaluation results show that SYNTHIFY outperforms the baseline in terms of both effectiveness and efficiency. We evaluated SYNTHIFY on 8 publicly-available control



systems and compared it with a state-of-the-art, open-source falsification tool PSY-TALiRo. These promising results demonstrate the high effectiveness and efficiency of SYNTHIFY in the falsification of AI-enabled control systems.

In the future, we plan to extend SYNTHIFY to AI-enabled control systems that take more complex inputs, such as images [120] and audio [8]. We also plan to integrate SYNTHIFY with other safety analysis techniques like runtime shields [85] and bug localization [112, 121] to further improve the efficiency of falsification in AI-enabled control systems. Furthermore, with the advent of Large Language Models (LLMs), investigating the use of code generation LLMs [86, 119, 122, 123] for synthesizing proxy programs is a promising avenue for further research.

**Replication Package.** For reproducibility, we have made our code and data publicly available: <https://github.com/soarsmu/Synthify>.

## ACKNOWLEDGMENTS

This research/project is supported by the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (AISG Award No: AISG2-RP-2020-017). This work was also supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP)-ICT Creative Consilience Program grant funded by the Korea government (MSIT) (IITP-2025-RS-2020-II201819). This work was also supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2021R1I1A3048013).

## REFERENCES

- [1] Houssam Abbas and Georgios Fainekos. 2012. Convergence proofs for Simulated Annealing falsification of safety properties. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. 1594–1601. <https://doi.org/10.1109/Allerton.2012.6483411>
- [2] Houssam Abbas, Andrew Winn, Georgios Fainekos, and A. Agung Julius. 2014. Functional gradient descent method for Metric Temporal Logic specifications. In *2014 American Control Conference*. 2312–2317. <https://doi.org/10.1109/ACC.2014.6859453>
- [3] Raja Ben Abdesslem, Annibale Panichella, Shiva Nejati, Lionel C. Briand, and Thomas Stifter. 2018. Testing autonomous cars for feature interaction failures using many-objective search. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (Montpellier, France) (ASE '18)*. Association for Computing Machinery, New York, NY, USA, 143–154. <https://doi.org/10.1145/3238147.3238192>
- [4] Arvind Adimoolam, Thao Dang, Alexandre Donzé, James Kapinski, and Xiaoqing Jin. 2017. Classification and Coverage-Based Falsification for Embedded Control Systems. In *Computer Aided Verification*, Rupak Majumdar and Viktor Kunčák (Eds.). Springer International Publishing, Cham, 483–503. [https://doi.org/10.1007/978-3-319-63387-9\\_24](https://doi.org/10.1007/978-3-319-63387-9_24)
- [5] Diego Aineto, Enrico Scala, Eva Onaindia, and Ivan Serina. 2023. Falsification of cyber-physical systems using PDDL+ planning. In *Proceedings of the Thirty-Third International Conference on Automated Planning and Scheduling (Prague, Czech Republic) (ICAPS '23)*. AAAI Press, Article 1, 5 pages. <https://doi.org/10.1609/icaps.v33i1.27172>
- [6] Yashwanth Annpureddy, Che Liu, Georgios Fainekos, and Sriram Sankaranarayanan. 2011. S-TaLiRo: A Tool for Temporal Logic Falsification for Hybrid Systems. In *Tools and Algorithms for the Construction and Analysis of Systems*, Parosh Aziz Abdulla and K. Rustan M. Leino (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 254–257. [https://doi.org/10.1007/978-3-642-19835-9\\_21](https://doi.org/10.1007/978-3-642-19835-9_21)
- [7] Andrea Arcuri and Lionel Briand. 2011. A practical guide for using statistical tests to assess randomized algorithms in software engineering. In *Proceedings of the 33rd International Conference on Software Engineering (Waikiki, Honolulu, HI, USA) (ICSE '11)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/1985793.1985795>
- [8] Muhammad Hilmi Asyrofi, Zhou Yang, Jieke Shi, Chu Wei Quan, and David Lo. 2021. Can Differential Testing Improve Automatic Speech Recognition Systems?. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 674–678. <https://doi.org/10.1109/ICSME52107.2021.00079>
- [9] Stanley Bak, Sergiy Bogomolov, Abdelrahman Hekal, Niklas Kochdumper, Ethan Lew, Andrew Mata, and Amir Rahmati. 2024. Falsification using Reachability of Surrogate Koopman Models. In *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control (Hong Kong SAR, China) (HSCC '24)*. Association for Computing Machinery, New York, NY, USA, Article 27, 13 pages. <https://doi.org/10.1145/3641513.3650141>

- [10] James Bergstra and Yoshua Bengio. 2012. Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research* 13, 10 (2012), 281–305. <http://jmlr.org/papers/v13/bergstra12a.html>
- [11] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. 2016. End to End Learning for Self-Driving Cars. arXiv:1604.07316 [cs.CV] <https://arxiv.org/abs/1604.07316>
- [12] Gary Brassington. 2017. Mean absolute error and root mean square error: which is the better metric for assessing model performance?. In *EGU General Assembly Conference Abstracts (EGU General Assembly Conference Abstracts)*. 3574. <https://ui.adsabs.harvard.edu/abs/2017EGUGA..19.3574B>
- [13] Dimo Brockhoff, Anne Auger, Nikolaus Hansen, Dirk V. Arnold, and Tim Hohm. 2010. Mirrored Sampling and Sequential Selection for Evolution Strategies. In *Parallel Problem Solving from Nature, PPSN XI*, Robert Schaefer, Carlos Cotta, Joanna Kolodziej, and Günter Rudolph (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 11–21. [https://doi.org/10.1007/978-3-642-15844-5\\_2](https://doi.org/10.1007/978-3-642-15844-5_2)
- [14] Runqi Chai, Antonios Tsourdos, Al Savvaris, Senchun Chai, Yuanqing Xia, and C.L. Philip Chen. 2021. Review of advanced guidance and control algorithms for space/aerospace vehicles. *Progress in Aerospace Sciences* 122 (2021), 100696. <https://doi.org/10.1016/j.paerosci.2021.100696>
- [15] Aniruddh Chandratte, Tomas Hernandez Acosta, Tanmay Khandait, Giulia Pedrielli, and Georgios Fainekos. 2023. Stealthy attacks formalized as STL formulas for Falsification of CPS Security. In *Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control* (San Antonio, TX, USA) (HSCC '23). Association for Computing Machinery, New York, NY, USA, Article 15, 8 pages. <https://doi.org/10.1145/3575870.3587122>
- [16] Pranav Singh Chib and Pravendra Singh. 2024. Recent Advancements in End-to-End Autonomous Driving Using Deep Learning: A Survey. *IEEE Transactions on Intelligent Vehicles* 9, 1 (2024), 103–118. <https://doi.org/10.1109/TIV.2023.3318070>
- [17] Edmund Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. 2000. Counterexample-Guided Abstraction Refinement. In *Computer Aided Verification*, E. Allen Emerson and Aravinda Prasad Sistla (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 154–169. [https://doi.org/10.1007/10722167\\_15](https://doi.org/10.1007/10722167_15)
- [18] Anthony Corso, Robert Moss, Mark Koren, Ritchie Lee, and Mykel Kochenderfer. 2022. A Survey of Algorithms for Black-Box Safety Validation of Cyber-Physical Systems. *J. Artif. Int. Res.* 72 (Jan. 2022), 377–428. <https://doi.org/10.1613/jair.1.12716>
- [19] Jonas Degraeve, Federico Felici, Jonas Buchli, Michael Neunert, Brendan Tracey, Francesco Carpanese, Timo Ewalds, Roland Hafner, Abbas Abdolmaleki, Diego de Las Casas, et al. 2022. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature* 602, 7897 (2022), 414–419. <https://doi.org/10.1038/s41586-021-04301-9>
- [20] Jyotirmoy Deshmukh, Marko Horvat, Xiaoqing Jin, Rupak Majumdar, and Vinayak S. Prabhu. 2017. Testing Cyber-Physical Systems through Bayesian Optimization. *ACM Trans. Embed. Comput. Syst.* 16, 5s, Article 170 (sep 2017), 18 pages. <https://doi.org/10.1145/3126521>
- [21] DJI. 2024. DJI Camera Drones. <https://www.dji.com/sg/camera-drones>
- [22] Alexandre Donzé. 2010. Breach, A Toolbox for Verification and Parameter Synthesis of Hybrid Systems. In *Computer Aided Verification*, Tayssir Touili, Byron Cook, and Paul Jackson (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 167–170. [https://doi.org/10.1007/978-3-642-14295-6\\_17](https://doi.org/10.1007/978-3-642-14295-6_17)
- [23] Alexandre Donzé and Oded Maler. 2010. Robust satisfaction of temporal logic over real-valued signals. In *Proceedings of the 8th International Conference on Formal Modeling and Analysis of Timed Systems (Klosterneuburg, Austria) (FORMATS'10)*. Springer-Verlag, Berlin, Heidelberg, 92–106. [https://doi.org/10.1007/978-3-642-15297-9\\_9](https://doi.org/10.1007/978-3-642-15297-9_9)
- [24] Tommaso Dreossi, Alexandre Donzé, and Sanjit A Seshia. 2019. Compositional falsification of cyber-physical systems with machine learning components. *JAR* 63, 4 (2019), 1031–1053. <https://doi.org/10.1007/s10817-018-09509-5>
- [25] Tommaso Dreossi, Daniel J. Fremont, Shromona Ghosh, Edward Kim, Hadi Ravanbakhsh, Marcell Vazquez-Chanlatte, and Sanjit A. Seshia. 2019. VerifAI: A Toolkit for the Formal Design and Analysis of Artificial Intelligence-Based Systems. In *Computer Aided Verification*, Isil Dillig and Serdar Tasiran (Eds.). Springer International Publishing, Cham, 432–442. [https://doi.org/10.1007/978-3-030-25540-4\\_25](https://doi.org/10.1007/978-3-030-25540-4_25)
- [26] Yan Duan, Xi Chen, Rein Houthoofd, John Schulman, and Pieter Abbeel. 2016. Benchmarking deep reinforcement learning for continuous control. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48* (New York, NY, USA) (ICML '16). JMLR.org, 1329–1338. <https://proceedings.mlr.press/v48/duan16.pdf>
- [27] Johan Lidén Eddeland, Alexandre Donzé, Sajed Miremadi, and Knut Åkesson. 2020. Industrial Temporal Logic Specifications for Falsification of Cyber-Physical Systems. 74 (2020), 267–274. <https://doi.org/10.29007/r74f>
- [28] Gidon Ernst, Paolo Arcaini, Ismail Bennani, Aniruddh Chandratte, Alexandre Donzé, Georgios Fainekos, Goran Frehse, Khoulood Gaaloul, Jun Inoue, Tanmay Khandait, Logan Mathesen, Claudio Menghi, Giulia Pedrielli, Marc Pouzet, Masaki Waga, Shakiba Yaghoubi, Yoriyuki Yamagata, and Zhenya Zhang. 2021. ARCH-COMP 2021 Category Report: Falsification with Validation of Results. In *8th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH21) (EPIc Series in Computing, Vol. 80)*, Goran Frehse and Matthias Althoff (Eds.). EasyChair, 133–152. <https://doi.org/10.29007/xw11>

- [29] Gidon Ernst, Paolo Arcaini, Ismail Bennani, Alexandre Donze, Georgios Fainekos, Goran Frehse, Logan Mathesen, Claudio Menghi, Giulia Pedrielli, Marc Pouzet, Shakiba Yaghoubi, Yoriyuki Yamagata, and Zhenya Zhang. 2020. ARCH-COMP 2020 Category Report: Falsification. *74* (2020), 140–152. <https://doi.org/10.29007/trr1>
- [30] Gidon Ernst, Paolo Arcaini, Georgios Fainekos, Federico Formica, Jun Inoue, Tanmay Khandait, Mohammad Mahdi Mahboob, Claudio Menghi, Giulia Pedrielli, Masaki Waga, Yoriyuki Yamagata, and Zhenya Zhang. 2022. ARCH-COMP 2022 Category Report: Falsification with Unbounded Resources. In *Proceedings of 9th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH22) (EPiC Series in Computing, Vol. 90)*, Goran Frehse, Matthias Althoff, Erwin Schoitsch, and Jeremie Guiochet (Eds.). EasyChair, 204–221. <https://doi.org/10.29007/fhmk>
- [31] Gidon Ernst, Sean Sedwards, Zhenya Zhang, and Ichiro Hasuo. 2021. Falsification of Hybrid Systems Using Adaptive Probabilistic Search. *ACM Trans. Model. Comput. Simul.* 31, 3, Article 18 (July 2021), 22 pages. <https://doi.org/10.1145/3459605>
- [32] Nathan Fulton, Nathan Hunt, Nghia Hoang, and Subhro Das. 2020. Formal Verification of End-to-End Learning in Cyber-Physical Systems: Progress and Challenges. (2020). arXiv:2006.09181 [cs.SE] <https://arxiv.org/abs/2006.09181>
- [33] Claire Glanois, Paul Weng, Matthieu Zimmer, Dong Li, Tianpei Yang, Jianye Hao, and Wulong Liu. 2024. A survey on interpretable reinforcement learning. *Machine Learning* 113 (2024), 5847–5890. <https://doi.org/10.1007/s10994-024-06543-w>
- [34] Divya Gopinath, Hayes Converse, Corina S. Păsăreanu, and Ankur Taly. 2020. Property inference for deep neural networks. In *Proceedings of the 34th IEEE/ACM International Conference on Automated Software Engineering (San Diego, California) (ASE '19)*. IEEE Press, 797–809. <https://doi.org/10.1109/ASE.2019.00079>
- [35] Huihui Guo, Fan Wu, Yunchuan Qin, Ruihui Li, Keqin Li, and Kenli Li. 2023. Recent Trends in Task and Motion Planning for Robotics: A Survey. *ACM Comput. Surv.* 55, 13s, Article 289 (July 2023), 36 pages. <https://doi.org/10.1145/3583136>
- [36] Fitash Ul Haq, Donghwan Shin, and Lionel Briand. 2022. Efficient online testing for DNN-enabled systems using surrogate-assisted and many-objective optimization. In *Proceedings of the 44th International Conference on Software Engineering (Pittsburgh, Pennsylvania) (ICSE '22)*. Association for Computing Machinery, New York, NY, USA, 811–822. <https://doi.org/10.1145/3510003.3510188>
- [37] Fitash Ul Haq, Donghwan Shin, and Lionel C. Briand. 2023. Many-Objective Reinforcement Learning for Online Testing of DNN-Enabled Systems. In *Proceedings of the 45th International Conference on Software Engineering (Melbourne, Victoria, Australia) (ICSE '23)*. IEEE Press, 1814–1826. <https://doi.org/10.1109/ICSE48619.2023.00155>
- [38] Fitash Ul Haq, Donghwan Shin, Shiva Nejati, and Lionel Briand. 2021. Can offline testing of deep neural networks replace their online testing? a case study of automated driving systems. *Empirical Software Engineering* 26, 5 (2021), 90. <https://doi.org/10.1007/s10664-021-09982-4>
- [39] Fitash Ul Haq, Donghwan Shin, Shiva Nejati, and Lionel C. Briand. 2020. Comparing Offline and Online Testing of Deep Neural Networks: An Autonomous Car Case Study. In *2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST)*. 85–95. <https://doi.org/10.1109/ICST46399.2020.00019>
- [40] Junda He, Zhou Yang, Jieke Shi, Chengran Yang, Kisub Kim, Bowen Xu, Xin Zhou, and David Lo. 2024. Curiosity-Driven Testing for Sequential Decision-Making Process. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 165, 14 pages. <https://doi.org/10.1145/3597503.3639149>
- [41] Todd Hester, Michael Quinlan, and Peter Stone. 2012. RTMBA: A Real-Time Model-Based Reinforcement Learning Architecture for robot control. In *2012 IEEE International Conference on Robotics and Automation*. IEEE, 85–90. <https://doi.org/10.1109/ICRA.2012.6225072>
- [42] Geoffrey E Hinton and Sam Roweis. 2002. Stochastic Neighbor Embedding. In *Advances in Neural Information Processing Systems*, S. Becker, S. Thrun, and K. Obermayer (Eds.), Vol. 15. MIT Press. [https://proceedings.neurips.cc/paper\\_files/paper/2002/file/6150ccc6069bea6b5716254057a194ef-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2002/file/6150ccc6069bea6b5716254057a194ef-Paper.pdf)
- [43] K.J. Hunt, D. Sbarbaro, R. Żbikowski, and P.J. Gawthrop. 1992. Neural networks for control systems—A survey. *Automatica* 28, 6 (1992), 1083–1112. [https://doi.org/10.1016/0005-1098\(92\)90053-1](https://doi.org/10.1016/0005-1098(92)90053-1)
- [44] Jeeyana Priya Inala, Osbert Bastani, Zenna Tavares, and Armando Solar-Lezama. 2020. Synthesizing Programmatic Policies that Inductively Generalize. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=S118oANFDH>
- [45] Ismailbennani. [n.d.]. zlscheck: A random testing tool for zélus. <https://github.com/ismailbennani/zlscheck>
- [46] Dario Izzo, Marcus Mörtens, and Binfeng Pan. 2019. A survey on artificial intelligence trends in spacecraft guidance dynamics and control. *Astrodynamics* 3 (2019), 287–299. <https://doi.org/10.1007/s42064-018-0053-6>
- [47] Michael A Johnson and Mohammad H Moradi. 2005. *PID control*. Springer, London. <https://doi.org/10.1007/1-84628-148-2>
- [48] Hao Ju, Rongshun Juan, Randy Gomez, Keisuke Nakamura, and Guangliang Li. 2022. Transferring policy of deep reinforcement learning from simulation to reality for robotics. *Nature Machine Intelligence* 4, 12 (2022), 1077–1087. <https://doi.org/10.1038/s42256-022-00573-6>
- [49] Kyle D. Julian, Jessica Lopez, Jeffrey S. Brush, Michael P. Owen, and Mykel J. Kochenderfer. 2016. Policy compression for aircraft collision avoidance systems. In *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*. 1–10. <https://doi.org/10.1109/DASC.2016.7778091>
- [50] M. M. Kamruzzaman. 2021. New Opportunities, Challenges, and Applications of Edge-AI for Connected Healthcare in Smart Cities. In *2021 IEEE Globecom Workshops (GC Wkshps)*. 1–6. <https://doi.org/10.1109/GCWkshps52748.2021.9682055>

- [51] Hong Jin Kang, Truong Giang Nguyen, Bach Le, Corina S. Păsăreanu, and David Lo. 2022. Test mimicry to assess the exploitability of library vulnerabilities. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis (Virtual, South Korea) (ISSTA 2022)*. Association for Computing Machinery, New York, NY, USA, 276–288. <https://doi.org/10.1145/3533767.3534398>
- [52] James Kapinski, Jyotirmoy V. Deshmukh, Xiaoqing Jin, Hisahiro Ito, and Ken Butts. 2016. Simulation-Based Approaches for Verification of Embedded Control Systems: An Overview of Traditional and Advanced Modeling, Testing, and Verification Techniques. *IEEE Control Systems Magazine* 36, 6 (2016), 45–64. <https://doi.org/10.1109/MCS.2016.2602089>
- [53] B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez. 2022. Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems* 23, 6 (2022), 4909–4926. <https://doi.org/10.1109/TITS.2021.3054625>
- [54] Amit Kishor and Chinmay Chakraborty. 2022. Artificial intelligence and internet of things based healthcare 4.0 monitoring system. *Wireless personal communications* 127, 2 (2022), 1615–1631. <https://doi.org/10.1007/s11277-021-08708-5>
- [55] Volodymyr Kuleshov and Doina Precup. 2014. Algorithms for multi-armed bandit problems. (2014). arXiv:1402.6028 [cs.AI] <https://arxiv.org/abs/1402.6028>
- [56] Atanu Kundu, Sauvik Gon, and Rajarshi Ray. 2024. Data-Driven Falsification of Cyber-Physical Systems. In *Proceedings of the 17th Innovations in Software Engineering Conference (Bangalore, India) (ISEC '24)*. Association for Computing Machinery, New York, NY, USA, Article 10, 5 pages. <https://doi.org/10.1145/3641399.3641401>
- [57] Tin-Yau Kwok and Dit-Yan Yeung. 1996. Use of bias term in projection pursuit learning improves approximation and convergence properties. *IEEE Transactions on Neural Networks* 7, 5 (1996), 1168–1183. <https://doi.org/10.1109/72.536312>
- [58] Matthew Landers and Afsaneh Doryab. 2023. Deep Reinforcement Learning Verification: A Survey. *ACM Comput. Surv.* 55, 14s, Article 330 (July 2023), 31 pages. <https://doi.org/10.1145/3596444>
- [59] Johan Lidén Eddeland, Alexandre Donzé, and Knut Åkesson. 2022. Multi-Requirement Testing Using Focused Falsification. In *Proceedings of the 25th ACM International Conference on Hybrid Systems: Computation and Control (Milan, Italy) (HSCC '22)*. Association for Computing Machinery, New York, NY, USA, Article 4, 11 pages. <https://doi.org/10.1145/3501710.3519521>
- [60] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. 2019. Continuous control with deep reinforcement learning. arXiv:1509.02971 [cs.LG] <https://arxiv.org/abs/1509.02971>
- [61] Guannan Lou, Yao Deng, Xi Zheng, Mengshi Zhang, and Tianyi Zhang. 2022. Testing of autonomous driving systems: where are we and where should we go?. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Singapore, Singapore) (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 31–43. <https://doi.org/10.1145/3540250.3549111>
- [62] Deyun Lyu, Jiayang Song, Zhenya Zhang, Zhijie Wang, Tianyi Zhang, Lei Ma, and Jianjun Zhao. 2023. AutoRepair: Automated Repair for AI-Enabled Cyber-Physical Systems under Safety-Critical Conditions. (2023). arXiv:2304.05617 [cs.SE] <https://arxiv.org/abs/2304.05617>
- [63] Oded Maler and Dejan Nickovic. 2004. Monitoring Temporal Properties of Continuous Signals. In *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Yassine Lakhnech and Sergio Yovine (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 152–166. [https://doi.org/10.1007/978-3-540-30206-3\\_12](https://doi.org/10.1007/978-3-540-30206-3_12)
- [64] Horia Mania, Aurelia Guy, and Benjamin Recht. 2018. Simple random search of static linear policies is competitive for reinforcement learning. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2018/file/7634ea65a4e6d9041cfd3f7de18e334a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/7634ea65a4e6d9041cfd3f7de18e334a-Paper.pdf)
- [65] H. B. Mann and D. R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics* 18, 1 (1947), 50 – 60. <https://doi.org/10.1214/aoms/1177730491>
- [66] Logan Mathesen, Giulia Pedrielli, and Georgios Fainekos. 2021. Efficient Optimization-Based Falsification of Cyber-Physical Systems with Multiple Conjunctive Requirements. In *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*. 732–737. <https://doi.org/10.1109/CASE49439.2021.9551474>
- [67] Ruijie Meng, Zhen Dong, Jialin Li, Ivan Beschastnikh, and Abhik Roychoudhury. 2022. Linear-time temporal logic guided greybox fuzzing. In *Proceedings of the 44th International Conference on Software Engineering (Pittsburgh, Pennsylvania) (ICSE '22)*. Association for Computing Machinery, New York, NY, USA, 1343–1355. <https://doi.org/10.1145/3510003.3510082>
- [68] Claudio Menghi, Shiva Nejati, Lionel Briand, and Yago Isasi Parache. 2020. Approximation-refinement testing of compute-intensive cyber-physical models: an approach based on system identification. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (Seoul, South Korea) (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 372–384. <https://doi.org/10.1145/3377811.3380370>
- [69] Robert J Moss, Mykel J Kochenderfer, Maxime Gariel, and Arthur Dubois. 2023. Bayesian Safety Validation for Black-Box Systems. In *AIAA AVIATION 2023 Forum*. 3596. <https://doi.org/10.2514/6.2023-3596>
- [70] Khan Muhammad, Amin Ullah, Jaime Lloret, Javier Del Ser, and Victor Hugo C. de Albuquerque. 2021. Deep Learning for Safe Autonomous Driving: Current Challenges and Future Directions. *IEEE Transactions on Intelligent Transportation Systems* 22, 7 (2021), 4316–4336. <https://doi.org/10.1109/TITS.2020.3032227>



- [71] Dejan Ničković and Tomoya Yamaguchi. 2020. RTAMT: Online Robustness Monitors from STL. In *Automated Technology for Verification and Analysis*, Dang Van Hung and Oleg Sokolsky (Eds.). Springer International Publishing, Cham, 564–571. [https://doi.org/10.1007/978-3-030-59152-6\\_34](https://doi.org/10.1007/978-3-030-59152-6_34)
- [72] NVIDIA. 2024. NVIDIA DRIVE. <https://developer.nvidia.com/drive>.
- [73] Mitchell Olsthoorn, Arie van Deursen, and Annibale Panichella. 2021. Generating highly-structured input data by combining search-based testing and grammar-based fuzzing. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering (Virtual Event, Australia) (ASE '20)*. Association for Computing Machinery, New York, NY, USA, 1224–1228. <https://doi.org/10.1145/3324884.3418930>
- [74] Fabio Palomba, Annibale Panichella, Andy Zaidman, Rocco Oliveto, and Andrea De Lucia. 2016. Automatic test case generation: what if test code quality matters?. In *Proceedings of the 25th International Symposium on Software Testing and Analysis (Saarbrücken, Germany) (ISSTA 2016)*. Association for Computing Machinery, New York, NY, USA, 130–141. <https://doi.org/10.1145/2931037.2931057>
- [75] Qi Pang, Yuanyuan Yuan, and Shuai Wang. 2022. MDPFuzz: testing models solving Markov decision processes. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis (Virtual, South Korea) (ISSTA 2022)*. Association for Computing Machinery, New York, NY, USA, 378–390. <https://doi.org/10.1145/3533767.3534388>
- [76] Jarkko Peltomäki and Ivan Porres. 2022. Falsification of Multiple Requirements for Cyber-Physical Systems Using Online Generative Adversarial Networks and Multi-Armed Bandits. In *2022 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. 21–28. <https://doi.org/10.1109/ICSTW55395.2022.00018>
- [77] Xin Qin, Yuan Xia, Aditya Zutshi, Chuchu Fan, and Jyotirmoy V. Deshmukh. 2022. Statistical Verification of Cyber-Physical Systems using Surrogate Models and Conformal Inference. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCP)*. 116–126. <https://doi.org/10.1109/ICCP54341.2022.00017>
- [78] Wenjie Qiu and He Zhu. 2022. Programmatic Reinforcement Learning without Oracles. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=6Tk2noBdvxt>
- [79] Aravind Rajeswaran, Kendall Lowrey, Emanuel V. Todorov, and Sham M Kakade. 2017. Towards Generalization and Simplicity in Continuous Control. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/9ddb9dd5d8aee9a76bf217a2a3c54833-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/9ddb9dd5d8aee9a76bf217a2a3c54833-Paper.pdf)
- [80] Zahra Ramezani, Kenan Šehić, Luigi Nardi, and Knut Åkesson. 2023. Falsification of Cyber-Physical Systems using Bayesian Optimization. (2023). [arXiv:2209.06735 \[eess.SY\]](https://arxiv.org/abs/2209.06735) <https://arxiv.org/abs/2209.06735>
- [81] Kristin Yvonne Rozier. 2016. Specification: The Biggest Bottleneck in Formal Methods and Autonomy. In *Verified Software. Theories, Tools, and Experiments*, Sandrine Blazy and Marsha Chechik (Eds.). Springer International Publishing, Cham, 8–26. [https://doi.org/10.1007/978-3-319-48869-1\\_2](https://doi.org/10.1007/978-3-319-48869-1_2)
- [82] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. (2017). [arXiv:1703.03864 \[stat.ML\]](https://arxiv.org/abs/1703.03864) <https://arxiv.org/abs/1703.03864>
- [83] Segway. 2024. Segway transporters. <https://ap-en.segway.com/>
- [84] Ridwan Shariffdeen, Yannic Noller, Lars Grunske, and Abhik Roychoudhury. 2021. Concolic program repair. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 390–405. <https://doi.org/10.1145/3453483.3454051>
- [85] Jieke Shi, Junda He, Zhou Yang, Đorđe Žikelić, and David Lo. 2024. Synthesizing Efficient and Permissive Programmatic Runtime Shields for Neural Policies. (2024). [arXiv:2410.05641 \[cs.SE\]](https://arxiv.org/abs/2410.05641) <https://arxiv.org/abs/2410.05641>
- [86] Jieke Shi, Zhou Yang, Junda He, Bowen Xu, and David Lo. 2022. Can Identifier Splitting Improve Open-Vocabulary Language Model of Code? . In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE Computer Society, Los Alamitos, CA, USA, 1134–1138. <https://doi.org/10.1109/SANER53432.2022.00130>
- [87] Jieke Shi, Zhou Yang, Hong Jin Kang, Bowen Xu, Junda He, and David Lo. 2024. Greening Large Language Models of Code. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Society (Lisbon, Portugal) (ICSE-SEIS'24)*. Association for Computing Machinery, New York, NY, USA, 142–153. <https://doi.org/10.1145/3639475.3640097>
- [88] Jieke Shi, Zhou Yang, and David Lo. 2024. Efficient and Green Large Language Models for Software Engineering: Vision and the Road Ahead. *ACM Trans. Softw. Eng. Methodol.* (Dec. 2024). <https://doi.org/10.1145/3708525> Just Accepted.
- [89] Jieke Shi, Zhou Yang, Bowen Xu, Hong Jin Kang, and David Lo. 2023. Compressing Pre-trained Models of Code into 3 MB. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (Rochester, MI, USA) (ASE '22)*. Association for Computing Machinery, New York, NY, USA, Article 24, 12 pages. <https://doi.org/10.1145/3551349.3556964>
- [90] Dominik Sobania, Dirk Schweim, and Franz Rothlauf. 2023. A Comprehensive Survey on Program Synthesis With Evolutionary Algorithms. *IEEE Transactions on Evolutionary Computation* 27, 1 (2023), 82–97. <https://doi.org/10.1109/TEVC.2022.3162324>
- [91] Armando Solar-Lezama. 2009. The Sketching Approach to Program Synthesis. In *Programming Languages and Systems*, Zhenjiang Hu (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 4–13. [https://doi.org/10.1007/978-3-642-10672-9\\_3](https://doi.org/10.1007/978-3-642-10672-9_3)

- [92] Mozhan Soltani, Annibale Panichella, and Arie van Deursen. 2020. Search-Based Crash Reproduction and Its Impact on Debugging. *IEEE Transactions on Software Engineering* 46, 12 (2020), 1294–1317. <https://doi.org/10.1109/TSE.2018.2877664>
- [93] Jiayang Song, Deyun Lyu, Zhenya Zhang, Zhijie Wang, Tianyi Zhang, and Lei Ma. 2022. When cyber-physical systems meet AI: a benchmark, an evaluation, and a way forward. In *Proceedings of the 44th International Conference on Software Engineering: Software Engineering in Practice (Pittsburgh, Pennsylvania) (ICSE-SEIP '22)*. Association for Computing Machinery, New York, NY, USA, 343–352. <https://doi.org/10.1145/3510457.3513049>
- [94] Jiayang Song, Xuan Xie, and Lei Ma. 2023. SIEGESIEGE: A Semantics-Guided Safety Enhancement Framework for AI-Enabled Cyber-Physical Systems. *IEEE Transactions on Software Engineering* 49, 8 (2023), 4058–4080. <https://doi.org/10.1109/TSE.2023.3282981>
- [95] Yang Sun, Christopher M. Poskitt, Jun Sun, Yuqi Chen, and Zijiang Yang. 2023. LawBreaker: An Approach for Specifying Traffic Laws and Fuzzing Autonomous Vehicles. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (Rochester, MI, USA) (ASE '22)*. Association for Computing Machinery, New York, NY, USA, Article 62, 12 pages. <https://doi.org/10.1145/3551349.3556897>
- [96] R.S. Sutton and A.G. Barto. 1998. Reinforcement Learning: An Introduction. *IEEE Transactions on Neural Networks* 9, 5 (1998), 1054–1054. <https://doi.org/10.1109/TNN.1998.712192>
- [97] Tesla. 2023. Autopilot. <https://www.thinkautonomous.ai/blog/tesla-end-to-end-deep-learning/>.
- [98] Quinn Thibeault, Jacob Anderson, Aniruddh Chandratre, Giulia Pedrielli, and Georgios Fainekos. 2021. PSY-TaLiRo: A Python Toolbox for Search-Based Test Generation for Cyber-Physical Systems. In *Formal Methods for Industrial Critical Systems*, Alberto Lluch Lafuente and Anastasia Mavridou (Eds.). Springer International Publishing, Cham, 223–231. [https://doi.org/10.1007/978-3-030-85248-1\\_15](https://doi.org/10.1007/978-3-030-85248-1_15)
- [99] Hoang-Dung Tran, Feiyang Cai, Manzananas Lopez Diego, Patrick Musau, Taylor T. Johnson, and Xenofon Koutsoukos. 2019. Safety Verification of Cyber-Physical Systems with Reinforcement Learning Control. *ACM Trans. Embed. Comput. Syst.* 18, 5s, Article 105 (Oct. 2019), 22 pages. <https://doi.org/10.1145/3358230>
- [100] Hoang-Dung Tran, Xiaodong Yang, Diego Manzananas Lopez, Patrick Musau, Luan Viet Nguyen, Weiming Xiang, Stanley Bak, and Taylor T Johnson. 2020. NNV: the neural network verification tool for deep neural networks and learning-enabled cyber-physical systems. In *International Conference on Computer Aided Verification*. Springer, 3–17. [https://doi.org/10.1007/978-3-030-53288-8\\_1](https://doi.org/10.1007/978-3-030-53288-8_1)
- [101] Dweep Trivedi, Jesse Zhang, Shao-Hua Sun, and Joseph J. Lim. 2024. Learning to synthesize programs as interpretable and generalizable policies. , Article 1926 (2024), 18 pages. <https://proceedings.neurips.cc/paper/2021/file/d37124c4c79f357cb02c655671a432fa-Paper.pdf>
- [102] Cumhuri Erkan Tuncali, Bardh Hoxha, Guohui Ding, Georgios Fainekos, and Sriram Sankaranarayanan. 2018. Experience Report: Application of Falsification Methods on the UxAS System. In *NASA Formal Methods*, Aaron Dutle, César Muñoz, and Anthony Narkawicz (Eds.). Springer International Publishing, Cham, 452–459. [https://doi.org/10.1007/978-3-319-77935-5\\_30](https://doi.org/10.1007/978-3-319-77935-5_30)
- [103] Udacity. 2024. Udacity Self-Driving Challenge #2: Predicting Steering Angles. <https://github.com/udacity/self-driving-car/tree/master/steering-models>
- [104] András Vargha and Harold D. Delaney. 2000. A Critique and Improvement of the “CL” Common Language Effect Size Statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (2000), 101–132. <http://www.jstor.org/stable/1165329>
- [105] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. 2018. Programmatically Interpretable Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, 5045–5054. <https://proceedings.mlr.press/v80/verma18a.html>
- [106] Kesav Viswanadha, Edward Kim, Francis Indaheng, Daniel J. Fremont, and Sanjit A. Seshia. 2021. Parallel and Multi-objective Falsification with Scenic and VeriAI. In *Runtime Verification: 21st International Conference, RV 2021, Virtual Event, October 11–14, 2021, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 265–276. [https://doi.org/10.1007/978-3-030-88494-9\\_15](https://doi.org/10.1007/978-3-030-88494-9_15)
- [107] George A. Vouros. 2022. Explainable Deep Reinforcement Learning: State of the Art and Challenges. *ACM Comput. Surv.* 55, 5, Article 92 (Dec. 2022), 39 pages. <https://doi.org/10.1145/3527448>
- [108] Masaki Waga. 2020. Falsification of cyber-physical systems with robustness-guided black-box checking. In *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control (Sydney, New South Wales, Australia) (HSCC '20)*. Association for Computing Machinery, New York, NY, USA, Article 11, 13 pages. <https://doi.org/10.1145/3365365.3382193>
- [109] Bohao Wang, Zhe Hou, Gelin Zhang, Jianqi Shi, and Yanhong Huang. 2021. Tree Ensemble Property Verification from A Testing Perspective. In *The 33rd International Conference on Software Engineering and Knowledge Engineering, SEKE 2021, KSIR Virtual Conference Center, USA, July 1 - July 10, 2021*, Shi-Kuo Chang (Ed.). KSI Research Inc., 166–171. <https://doi.org/10.18293/SEKE2021-087>
- [110] Renzhi Wang, Zhehua Zhou, Jiayang Song, Xuan Xie, Xiaofei Xie, and Lei Ma. 2024. MORTAR: A Model-based Runtime Action Repair Framework for AI-enabled Cyber-Physical Systems. (2024). arXiv:2408.03892 [cs.SE] <https://arxiv.org/abs/2408.03892>
- [111] Xu Wang, Sen Wang, Xingxing Liang, Dawei Zhao, Jincai Huang, Xin Xu, Bin Dai, and Qiguang Miao. 2024. Deep Reinforcement Learning: A Survey. *IEEE Transactions on Neural Networks and Learning Systems* 35, 4 (2024), 5064–5078. <https://doi.org/10.1109/TNNLS.2022.3207346>
- [112] Ratnadira Widayarsi, Stefanus Agus Haryono, Ferdian Thung, Jieke Shi, Constance Tan, Fiona Wee, Jack Phan, and David Lo. 2022. On the Influence of Biases in Bug Localization: Evaluation and Benchmark. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 128–139. <https://doi.org/10.1109/SANER53432.2022.00027>



- [113] Yuanqing Wu, Siqin Liao, Xiang Liu, Zhihang Li, and Renquan Lu. 2023. Deep Reinforcement Learning on Autonomous Driving Policy With Auxiliary Critic Network. *IEEE Transactions on Neural Networks and Learning Systems* 34, 7 (2023), 3680–3690. <https://doi.org/10.1109/TNNLS.2021.3116063>
- [114] Yuan-Qing Xia, Yu-Long Gao, Li-Ping Yan, and Meng-Yin Fu. 2015. Recent progress in networked control systems—A survey. *International Journal of Automation and Computing* 12, 4 (2015), 343–367. <https://doi.org/10.1007/s11633-015-0894-x>
- [115] Xuan Xie, Jiayang Song, Zehua Zhou, Fuyuan Zhang, and Lei Ma. 2023. Mosaic: Model-based Safety Analysis Framework for AI-enabled Cyber-Physical Systems. (2023). arXiv:2305.03882 [cs.SE] <https://arxiv.org/abs/2305.03882>
- [116] Shakiba Yaghoubi and Georgios Fainekos. 2017. Hybrid approximate gradient and stochastic descent for falsification of nonlinear systems. In *2017 American Control Conference (ACC)*. 529–534. <https://doi.org/10.23919/ACC.2017.7963007>
- [117] Shakiba Yaghoubi and Georgios Fainekos. 2019. Gray-box adversarial testing for control systems with machine learning components. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control (Montreal, Quebec, Canada) (HSCC '19)*. Association for Computing Machinery, New York, NY, USA, 179–184. <https://doi.org/10.1145/3302504.3311814>
- [118] Yoriyuki Yamagata, Shuang Liu, Takumi Akazaki, Yihai Duan, and Jianye Hao. 2021. Falsification of Cyber-Physical Systems Using Deep Reinforcement Learning. *IEEE Transactions on Software Engineering* 47, 12 (2021), 2823–2840. <https://doi.org/10.1109/TSE.2020.2969178>
- [119] Chengran Yang, Hong Jin Kang, Jieke Shi, and David Lo. 2024. ACECode: A Reinforcement Learning Framework for Aligning Code Efficiency and Correctness in Code Language Models. (2024). arXiv:2412.17264 [cs.SE] <https://arxiv.org/abs/2412.17264>
- [120] Zhou Yang, Jieke Shi, Muhammad Hilmi Asyrofi, and David Lo. 2022. Revisiting Neuron Coverage Metrics and Quality of Deep Neural Networks. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 408–419. <https://doi.org/10.1109/SANER53432.2022.00056>
- [121] Zhou Yang, Jieke Shi, Shaowei Wang, and David Lo. 2021. IncBL: Incremental Bug Localization. In *2021 36th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 1223–1226. <https://doi.org/10.1109/ASE51524.2021.9678546>
- [122] Zhou Yang, Zhipeng Zhao, Chenyu Wang, Jieke Shi, Dongsun Kim, DongGyun Han, and David Lo. 2024. Gotcha! This Model Uses My Code! Evaluating Membership Leakage Risks in Code Models. *IEEE Transactions on Software Engineering* 50, 12 (2024), 3290–3306. <https://doi.org/10.1109/TSE.2024.3482719>
- [123] Zhou Yang, Zhipeng Zhao, Chenyu Wang, Jieke Shi, Dongsun Kim, Donggyun Han, and David Lo. 2024. Unveiling Memorization in Code Models. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 72, 13 pages. <https://doi.org/10.1145/3597503.3639074>
- [124] Dan Zhang, Peng Shi, Qing-Guo Wang, and Li Yu. 2017. Analysis and synthesis of networked control systems: A survey of recent advances and challenges. *ISA Transactions* 66 (2017), 376–392. <https://doi.org/10.1016/j.isatra.2016.09.026>
- [125] Shaohua Zhang, Shuang Liu, Jun Sun, Yuqi Chen, Wenzhi Huang, Jinyi Liu, Jian Liu, and Jianye Hao. 2022. FIGCPS: effective failure-inducing input generation for cyber-physical systems with deep reinforcement learning. In *Proceedings of the 36th IEEE/ACM International Conference on Automated Software Engineering (Melbourne, Australia) (ASE '21)*. IEEE Press, 555–567. <https://doi.org/10.1109/ASE51524.2021.9678832>
- [126] Xiaodong Zhang, Wei Zhao, Yang Sun, Jun Sun, Yulong Shen, Xuewen Dong, and Zijiang Yang. 2023. Testing Automated Driving Systems by Breaking Many Laws Efficiently. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis (Seattle, WA, USA) (ISSTA 2023)*. Association for Computing Machinery, New York, NY, USA, 942–953. <https://doi.org/10.1145/3597926.3598108>
- [127] Zhenya Zhang and Paolo Arcaini. 2021. Gaussian Process-Based Confidence Estimation for Hybrid System Falsification. In *Formal Methods, Marieke Huisman, Corina Păsăreanu, and Naijun Zhan (Eds.)*. Springer International Publishing, Cham, 330–348. [https://doi.org/10.1007/978-3-030-90870-6\\_18](https://doi.org/10.1007/978-3-030-90870-6_18)
- [128] Zhenya Zhang, Gidon Ernst, Sean Sedwards, Paolo Arcaini, and Ichiro Hasuo. 2018. Two-Layered Falsification of Hybrid Systems Guided by Monte Carlo Tree Search. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 37, 11 (2018), 2894–2905. <https://doi.org/10.1109/TCAD.2018.2858463>
- [129] Zhenya Zhang, Deyun Lyu, Paolo Arcaini, Lei Ma, Ichiro Hasuo, and Jianjun Zhao. 2021. Effective Hybrid System Falsification Using Monte Carlo Tree Search Guided by QB-Robustness. In *Computer Aided Verification, Alexandra Silva and K. Rustan M. Leino (Eds.)*. Springer International Publishing, Cham, 595–618. [https://doi.org/10.1007/978-3-030-81685-8\\_29](https://doi.org/10.1007/978-3-030-81685-8_29)
- [130] Zhenya Zhang, Deyun Lyu, Paolo Arcaini, Lei Ma, Ichiro Hasuo, and Jianjun Zhao. 2021. On the Effectiveness of Signal Rescaling in Hybrid System Falsification. In *NASA Formal Methods: 13th International Symposium, NFM 2021, Virtual Event, May 24–28, 2021, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 392–399. [https://doi.org/10.1007/978-3-030-76384-8\\_24](https://doi.org/10.1007/978-3-030-76384-8_24)
- [131] Zhenya Zhang, Deyun Lyu, Paolo Arcaini, Lei Ma, Ichiro Hasuo, and Jianjun Zhao. 2023. FalsifAI: Falsification of AI-Enabled Hybrid Control Systems Guided by Time-Aware Coverage Criteria. *IEEE Trans. Softw. Eng.* 49, 4 (April 2023), 1842–1859. <https://doi.org/10.1109/TSE.2022.3194640>
- [132] Zehua Zhou, Jiayang Song, Xuan Xie, Zhan Shu, Lei Ma, Dikai Liu, Jianxiong Yin, and Simon See. 2024. Towards Building AI-CPS with NVIDIA Isaac Sim: An Industrial Benchmark and Case Study for Robotics Manipulation. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice (Lisbon, Portugal) (ICSE-SEIP '24)*. Association for Computing

- Machinery, New York, NY, USA, 263–274. <https://doi.org/10.1145/3639477.3639740>
- [133] Zhehua Zhou, Xuan Xie, Jiayang Song, Zhan Shu, and Lei Ma. 2024. GenSafe: A Generalizable Safety Enhancer for Safe Reinforcement Learning Algorithms Based on Reduced Order Markov Decision Process Model. *IEEE Transactions on Neural Networks and Learning Systems* (2024), 1–15. <https://doi.org/10.1109/TNNLS.2024.3496492>
- [134] He Zhu, Zikang Xiong, Stephen Magill, and Suresh Jagannathan. 2019. An inductive synthesis framework for verifiable reinforcement learning. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation* (Phoenix, AZ, USA) (*PLDI 2019*). Association for Computing Machinery, New York, NY, USA, 686–701. <https://doi.org/10.1145/3314221.3314638>
- [135] Amirhossein Zolfagharian, Manel Abdellatif, Lionel C. Briand, Mojtaba Bagherzadeh, and Ramesh S. 2023. A Search-Based Testing Approach for Deep Reinforcement Learning Agents . *IEEE Transactions on Software Engineering* 49, 07 (July 2023), 3715–3735. <https://doi.org/10.1109/TSE.2023.3269804>