

Detecting and Mitigating System-Level Anomalies of Vision-Based Controllers

Aryaman Gupta^{*1}, Kaustav Chakraborty^{*2}, Somil Bansal²

Abstract—Autonomous systems, such as self-driving cars and drones, have made significant strides in recent years by leveraging visual inputs and machine learning for decision-making and control. Despite their impressive performance, these vision-based controllers can make erroneous predictions when faced with novel or out-of-distribution inputs. Such errors can cascade to catastrophic system failures and compromise system safety. In this work, we introduce a run-time anomaly monitor to detect and mitigate such closed-loop, system-level failures. Specifically, we leverage a reachability-based framework to stress-test the vision-based controller offline and mine its system-level failures. This data is then used to train a classifier that is leveraged online to flag inputs that might cause system breakdowns. The anomaly detector highlights issues that transcend individual modules and pertain to the safety of the overall system. We also design a fallback controller that robustly handles these detected anomalies to preserve system safety. We validate the proposed approach on an autonomous aircraft taxiing system that uses a vision-based controller for taxiing. Our results show the efficacy of the proposed approach in identifying and handling system-level anomalies, outperforming methods such as prediction error-based detection, and ensembling, thereby enhancing the overall safety and robustness of autonomous systems.

Website: phoenixrider12.github.io/FailureMitigation

I. INTRODUCTION

With the advances in deep learning and computer vision, modern autonomous and robotic systems have reached a level of competence that, in some instances, exceeds human capabilities [1]. Nevertheless, given the vast array of scenarios these vision-driven controllers might face in the real world, we can never entirely rule out the occurrence of uncommon corner cases and failure scenarios. Thus, even as we aspire for our robots to adapt to new conditions, there is a growing need for runtime anomaly detection systems that can provide early alerts when a system encounters anomalies, helping to counteract and mitigate potential rare failures [2].

Anomaly detection (AD) methods for learning components typically fall under two categories: distributional shift methods and functional uncertainty methods [3]. The former aims to detect and mitigate distribution shifts between training and the test times. These methods include approaches that artificially inject noise in the ground truth data (e.g., images) to force a distribution shift and create anomalous data [4].

^{*}The first two authors contributed equally to this paper.

¹Author is with the ECE Department at the Indian Institute of Technology (IIT) Varanasi. aryaman.gupta.ece20@iitbhu.ac.in.

²Authors are with the ECE Department at the University of Southern California. [kaustavc, somilban](mailto:{kaustavc, somilban}@usc.edu)@usc.edu. This research is supported in part by the DARPA Assured Neuro Symbolic Learning and Reasoning (ANSR) program and by the NSF CAREER program (2240163). Aryaman was supported by the IUSSTF-Viterbi Summer Research Program.




Fig. 1: Detecting and mitigating system-level anomalies of vision based controller. *Offline* (top row): We compute the Backward Reachable Tube (BRT) of the vision-based system using Hamilton-Jacobi Reachability analysis [19]. This allows us to create a labeled dataset of safe and unsafe images without any manual intervention. We train the anomaly detector using this dataset. *Online* (bottom row): We run the system on previously unseen environments where the trained anomaly detector triggers a fallback controller or the default controller depending on whether the observed image is anomalous.

Another line of research in this direction aims to develop learning algorithms that are distributionally robust, optimizing the worst-case performance within a pre-specified envelope of distributional shifts to guarantee out-of-distribution performance [5], [6] or performing domain randomization during training [7]. Since detecting distribution shifts can be challenging in general, especially when the test distribution is unknown *a priori*, the functional uncertainty methods instead detect the inputs that are either dissimilar to the training data [8]–[12] or lead to erroneous or low-confidence predictions [13]–[18]. The above methods predominately detect anomalous inputs at the component level; however, such component level monitoring (e.g., detecting image classification errors) can be insufficient to prevent system-level faults. Seemingly minor errors in the individual modules can cascade into catastrophic effects at a system level. Hence, a system-level view of such problems is often encouraged and is the main motivation behind this work.

In this work, we present an approach to detect and mitigate such system-level anomalies for autonomous systems that leverage learning-driven vision-based controllers for decision making (Fig. 1). Our work builds upon the study conducted in [19] that leverages a visual simulator within a Hamilton-Jacobi Reachability framework to expose the closed-loop failures of the system under the vision-based controller. However, the proposed framework requires privileged information about the test environment and is computationally not suitable for online applications. Our key idea is to utilize

this framework *offline* to stress-test and automatically mine the system-level failures of the closed-loop system across a variety of environment conditions. Ultimately, this process provides us with a diverse set of input image sequences which when seen by the vision-based controller leads to an overall system failure. These images are then used to train a simple classifier that can serve as an anomaly detector during runtime. We demonstrate that the resultant AD implicitly leverages the revealed failure modes to classify whether a previously unseen input image has the potential to trigger a system failure, without requiring any domain-specific heuristics. Next, we employ the AD to trigger a simple fallback controller online that can trade the system's performance to assure system safety whenever the anomaly detector determines the system is at risk of entering the failure zone. We illustrate the joint properties of our anomaly detector and fallback controller on an autonomous aircraft taxiing case study, leveraging a vision-based controller. We compare our method against commonly used component-level AD techniques, such as predicting ensemble-based uncertainty, and prediction error-based anomaly detectors to highlight the key advantages of the proposed system-level anomaly detector.

II. PROBLEM SETUP

Consider a robot with dynamics, $\dot{\mathbf{x}} = f(\mathbf{x}, u)$ where, state $\mathbf{x} \in \mathbb{R}^n$ and control $u \in \mathbb{U}$ (a compact set). The robot possesses a sensor \mathcal{S} , that allows it to perceive visual inputs from its surroundings at any given state \mathbf{x} . We can express the mapping between the state to the input as, $I = \mathcal{S}(\mathbf{x})$. I could be an RGB image, a pointcloud, etc. Additionally, we have a vision-based controller, π , that maps I to the control u , and defined as, $u := \pi(I)$. Note that π can be an end-to-end policy or it can consist of several sub-modules some of which may not be data-driven.

Let $\zeta_{\mathbf{x}}^{\pi}(\tau)$ be the robot's state achieved at time τ when it starts from state \mathbf{x} at time $t = 0$, and follows the policy π over $[0, \tau]$. Finally, let \mathcal{O} denote a set of undesirable states (or failure states) for the system. As an example, \mathcal{O} could represent obstacles for a ground robot. Thus, an initial state \mathbf{x} is considered unsafe for the system if $\exists s \in [0, \tau], \zeta_{\mathbf{x}}^{\pi}(s) \in \mathcal{O}$. The set of input images the system sees, starting from such unsafe states, are thus considered *anomalous for the closed-loop system* (as they eventually steer the system to \mathcal{O}) and denoted as \mathcal{I}_{unsafe} .

In this work, we are primarily interested in obtaining a mapping, σ , that provides a binary decision of whether a given input I can possibly lead to the failure of the system:

$$\sigma : I \rightarrow \{0, 1\} \quad (1)$$

where 1 means that I is anomalous, and 0 means it is not. Hence, an ideal σ should output 1 whenever $I \in \mathcal{I}_{unsafe}$ and 0 otherwise. In addition, we seek to find a mitigation system that can preserve system safety even if the robot encounters an anomalous input in \mathcal{I}_{unsafe} .

Generating σ requires addressing a few challenges: (a) \mathcal{I}_{unsafe} is typically hard to obtain as the test environment is

unknown *a priori*; (b) even in known environments, obtaining high-dimensional inputs, e.g., RGB images, that lead to system-level failures is a challenging problem; finally, (c) for vision-based controllers, \mathcal{I}_{unsafe} varies with changes in the robot's surroundings, e.g., \mathcal{I}_{unsafe} for an indoor navigation robot might change as we move from one room to another. Thus, formulation of a general AD pipeline that can recognize features of anomalous input images without prior knowledge of the new surroundings is non-trivial.

Running example (TaxiNet). We introduce the autonomous aircraft taxiing problem [20] as a running example to illustrate the key aspects of our framework. Here, the robot is a Cessna 208B Grand Caravan aircraft modeled as a three-dimensional non-linear system with dynamics:

$$\dot{p}_x = v \cos(\theta) \quad \dot{p}_y = v \sin(\theta) \quad \dot{\theta} = u \quad (2)$$

where p_x is the crosstrack error (CTE), p_y is the downtrack position (DTP) and θ is the heading error (HE) of the aircraft in degrees from the centreline (Fig. 2(a) shows how these quantities are measured). v is the linear velocity of the aircraft kept constant at 5 m/s, and the control u is the angular velocity.




Fig. 2: (a) p_x, p_y, θ denote the state of the aircraft; dashed-white lines show FoV of the camera. Runway simulation images with clear sky at (b) 9AM, (c) 5PM, and (d) 9PM, and overcast clouds at (e) 9AM, (f) 5PM, and (g) 9PM taken by camera mounted on the aircraft for the KMWH runway showing the variations in lighting conditions and shadows ((c),(f)) for changes in the environment.

The goal of the aircraft is to follow the centreline as closely as possible using the images obtained through a camera mounted on its right wing. For this purpose, the aircraft uses a Convolutional Neural Network (CNN), which returns the estimated CTE and HE, $(\hat{p}_x, \hat{\theta})$. A proportional controller (P-Controller) then takes these predicted tracking errors to return the control input as follows:

$$u := \tan(-0.74\hat{p}_x - 0.44\hat{\theta}) \quad (3)$$

Hence, the policy π is a composition of the CNN and the P-Controller. Intuitively, the P-controller is designed to steer the aircraft towards the centreline based on the state estimate provided by the CNN [20]. We are interested in analyzing this *given* controller without modifying its underlying design.

The image observations are obtained using the X-Plane flight simulator that can render the RGB image, I , from a virtual camera (\mathcal{S}) mounted on the right wing of the aircraft at any state and a given time of day (see Fig. 2(b)-(g) for representative images under different simulation conditions).

We define the unsafe states for the aircraft as $\mathcal{O} = \{\mathbf{x} : |p_x| \geq B\}$, where B is the runway width. Thus, \mathcal{O} corresponds to aircraft leaving the runway. Our goal is to find the mapping, σ , that is able to classify whether an input image I is likely to eventually drive the system off the runway.

III. BACKGROUND: HAMILTON-JACOBI REACHABILITY

In this work, we will use Hamilton-Jacobi (HJ) Reachability analysis to obtain a dataset of system-level anomalies offline. We now provide a brief overview of HJ reachability and refer the readers to [21] for more details.

In reachability analysis, we focus on calculating the *Backward Reachable Tube (BRT)* of the system. The BRT refers to the collection of initial states from which an agent following policy $\pi(\mathbf{x})$, can reach the target set \mathcal{O} within the time interval $[t, T]$:

$$\mathcal{V} := \{\mathbf{x} : \exists \tau \in [t, T], \zeta_{\mathbf{x}}^{\pi}(\tau) \in \mathcal{O}\} \quad (4)$$

HJ reachability analysis allows us to compute the BRT for general nonlinear systems. To compute the BRT, the target set is first represented as a sub-zero level set of a function $l(\mathbf{x})$, denoted as $\mathcal{O} = \{\mathbf{x} : l(\mathbf{x}) \leq 0\}$ [22], [23]. The function $l(\mathbf{x})$ typically represents the signed distance from a state to the target set \mathcal{O} . With this formulation, the BRT computation can be reframed as an optimal control problem that involves finding a value function defined as:

$$V(\mathbf{x}, t) = \min_{\tau \in [t, T]} l(\zeta_{\mathbf{x}}^{\pi}(\tau)) \quad (5)$$

This value function (defined in (5)), can be iteratively computed using dynamic programming principles leading to a partial differential equation known as the Hamilton-Jacobi-Bellman Variational Inequality (HJB-VI) [21]:

$$\begin{aligned} \min\{D_t V(\mathbf{x}, t) + H(\mathbf{x}, t), l(\mathbf{x}) - V(\mathbf{x}, t)\} = 0 \\ \text{with } V(\mathbf{x}, T) = l(\mathbf{x}) \end{aligned} \quad (6)$$

In this equation, D_t and ∇ represent the time and spatial gradients of the value function, respectively. The Hamiltonian, denoted as $H := \langle \nabla V(\mathbf{x}, t), f(\mathbf{x}, \pi(\mathbf{x})) \rangle$, embeds the system dynamics in the HJB-VI. Consequently, the BRT corresponds to the subzero level set of the value function:

$$\mathcal{V} = \{\mathbf{x} : V(\mathbf{x}, t) \leq 0\} \quad (7)$$

In the next section, we describe how we can use \mathcal{V} for obtaining a system-level anomaly detector.

IV. LEARNING AND MITIGATING SYSTEM FAILURES USING REACHABILITY ANALYSIS

In this work, we use a learned classifier model as our anomaly detector. Our key idea is to compute the BRT of the system offline under a vision-based controller for a diverse set of environments. The BRT can then be used to label the training data for our classifier. We now describe this process in detail over three steps.

A. Automatic labels for system-level anomalous inputs

To compute a dataset of anomalous inputs, we first compute the set of all starting states that lead the closed-loop system to \mathcal{O} under the vision-based controller. In other words, we compute the BRT of the system under the vision-based controller. Once the BRT is obtained, the images corresponding to the states inside the BRT can be used as examples of system-level anomalous inputs, i.e.,

$$I \in \mathcal{I}_{unsafe} \Leftrightarrow \mathbf{x} \in \mathcal{V}, \text{ where, } I = S(\mathbf{x}) \quad (8)$$

However, the BRT computation in (6) typically requires an analytical model of $\pi(x)$, which is not possible in our case,

since an analytical model of vision sensor S is generally not available. To overcome this challenge, we will follow the approach used in [19], wherein the BRT is computed using image samples collected from a photorealistic simulator. These image samples are then processed through the vision-based controller to obtain samples of u , which are subsequently used to approximate the BRT. Thus, access to photorealistic simulators provides us with an opportunity to inexpensively sample from high-dimensional state spaces, which could otherwise be a challenging and tedious process. Moreover, since the image observed at a particular state depends on the environment conditions (e.g., lighting conditions, weather, etc.), we compute BRTs for a diverse set of environment conditions to capture a wider set of system anomalies in our training dataset, allowing for a better generalization during runtime.

Once all the BRTs are computed, we can sample states randomly in the state space, render the corresponding images, and automatically label any sampled image as safe or anomalous depending on whether the state is outside or inside the BRT. Note that by the construction of the training dataset, our method targets image inputs that lead to system-level failures, without requiring any manual labeling.

B. Learning an anomaly detector

Equipped with a training dataset of system-level anomalies, we train a binary classifier to predict the label for a novel input image. Specifically, we use a deep neural network (DNN)-based classifier, given their remarkable success with image classification tasks. The DNN takes as input an image and returns the softmax scores of the input being an anomaly. During training, we assume access to the system's BRT and hence the true labels of an image being anomalous. However, during testing, we do not possess the BRT information. Our model is seen to identify anomalies during testing without needing the BRT by learning the underlying features of anomalous inputs. We evaluate the performance and highlight some interesting cases of the detected anomalies by the trained classifier in Sec. V.

C. Fallback controller

Equipped with an anomaly detector, we propose a safety-preserving controller pipeline for the system. Specifically, when the system detects a possible anomalous input, we switch the system's vision-based controller π with a fallback controller π^* . For an input I , the overall control input for the system is given as:

$$u_{\text{filtered}} = \begin{cases} \pi^*(I), & \text{if } \sigma(I) = 1 \\ \pi(I), & \text{otherwise,} \end{cases} \quad (9)$$

where σ is the learned anomaly detector. The condition $\sigma(I) = 1$ is triggered if I is an anomaly, and hence the system switches to π^* .

The choice of π^* is primarily aimed at maintaining system safety and may vary from system to system. For example, π^* might be an alternative controller that is safer but computationally intensive (e.g., a full-stack SLAM for navigation) or

trade-off performance for safety (e.g., coming to a complete stop), or logistically expensive (e.g., relying on a human operator). Designing a good π^* for a system is itself an interesting research question that we defer to future work.

V. CASE STUDY: AUTONOMOUS AIRCRAFT TAXIING

Generation of the Anomaly Dataset. To generate data for training the AD, we compute BRT of the aircraft under the TaxiNet controller across two different times of the day (9 AM and 9 PM), two cloud conditions (clear and overcast), and three different runways (codenamed: KMWH, KATL, PAEI), leading to 12 different BRTs. Note that each of the above conditions might lead to a different visual input I at the same state \mathbf{x} (see Fig. 2 (b)-(e) for samples images), and hence lead to a different control input being applied to the aircraft, and subsequently, different BRTs.

To compute the BRTs, we use the Level Set Toolbox (LST) [24] that solves the HJB-VI in (6) numerically over a state-space grid for a time horizon of 8s. Following [19], we use a uniform $101 \times 101 \times 101$ grid over $p_x \in [-X, X]m$, $p_y \in [100, 250]m$ and $\theta \in [-28^\circ, 28^\circ]$ (the value of X depends on the runway, e.g., for KMWH, $X = 11$). The LST requires the control input u at each of the grid point to compute the BRT, which is obtained by rendering the image using X-Plane simulator at that state and querying the controller in (3). This consists of passing the rendered image through the TaxiNet CNN, followed by a P-controller. We refer the readers to [19] for more details on the BRT computation.

Fig. 3 (left) shows a slice of the computed BRT for the KMWH runway at 9am in clear conditions for $p_y = 100m$. The gray area represent the set of starting states of the aircraft from which it will eventually leave the runway under TaxiNet, whereas white represents the safe area. We next randomly sample 20K states for each of the 12 conditions and render the images for these states, leading to an overall training dataset of 240K images. If an image is generated from a state present in the BRT, we label it as an anomaly. Otherwise, we label it as safe.




Fig. 3: BRT of 9AM KMWH runway (part of train set) shown on the left along with a few training images and BRT of 5PM KSFO runway (part of test set) shown on the right, along with a few testing images, showing diversity in our training and testing scenarios.

Anomaly Detector. We next train a binary classifier on the collected dataset that takes a $224 \times 224 \times 3$ input image and returns the probability of the image being anomalous. We use a pre-trained EfficientNet-B0 Model [25] as our backbone, replacing the last layer with a fully connected layer that feeds into a binary softmax output layer. We trained the classifier using cross-entropy loss and Adam optimizer for 20 epochs

with a fixed learning rate of $3e^{-4}$ which took around 1 hour on a Nvidia RTX 3090 GPU.

Evaluation. To test the generalization capabilities of the learned AD, we test it on an unseen time of the day (5 PM) for the three airports present in the training set, as well as on two unseen KSFO and KEWR airports, across both cloud conditions (clear and overcast). Specifically, we measure the recall and accuracy of the learned AD. A higher recall value tells us that our AD can reliably detect true positives, i.e., the AD returns a conservative prediction if it is unsure of an input. Such a behavior err on caution to produce a safety-first detector. To compute these metrics, we also obtain the BRT of the system for the test environments; however, note that during testing, our network has access to only the sampled image *without* any privileged information regarding the system, surroundings, or the BRTs. We sum-

TABLE I: Performance of the learned anomaly detector.

Airport ID	Recall (%)		Accuracy (%)	
	(C)	(O)	(C)	(O)
KMWH	91.50	92.22	95.89	97.02
KATL	97.44	95.89	98.19	97.97
PAEI	93.24	93.71	97.71	97.99
KEWR	96.31	94.20	96.29	96.72
KSFO	90.34	90.56	89.72	90.62

marize the performance of the anomaly detector in Table I (Here (C) = cloudy, (O) = overcast). As evident from the results, the learned AD is consistently able to detect system-level anomalies in new environments. We illustrate some of the representative images that were classified as anomalous in Fig. 4. The AD is able to learn the safe limits of the runway and that if the system is close to the runway boundary. TaxiNet may not be able to correctly estimate aircraft's state and it might fail under the vision-based controller. This may not be particularly surprising as most states near the runway boundary (white lines in Fig. 4 (a-b)) are unsafe in the training dataset (the region near the runway boundary is contained in the BRT in Fig. 3), potentially allowing the AD to learn that such images often lead to a failure. However, what's interesting is that the AD learns that a similar image *will not* cause a system failure during night time because of the runway lights, which when are lit, help the TaxiNet to estimate the aircraft position more accurately, avoiding a failure. Finally, we noticed that several images classified as anomalies contained the aircraft runway markings (Fig. 4 (d-f)). Such an emergent understanding of semantic failure mode is quite impressive since in the prior work [19], the authors had to manually analyze the failures to recognize these patterns as anomalies. Our AD justifies the use of a learned classifier by practically automating the detection process without any inclusion of heuristics or manual intervention.

Baselines. We next compare our method to a few commonly



Fig. 4: Failures detected by AD. (a, b) Images correspond to the aircraft being close to the runway boundaries (highlighted with the magenta bounding boxes). (c, d) The visual controller confuses the runway markings (highlighted with the cyan bounding boxes) with the centerline and ultimately leads to a system failure. (e, f) Image (f) is (accurately) not classified as an anomaly during the night time (the same image is classified as anomaly during the day, shown in (e)), as the runway lights (highlighted with the yellow bounding boxes) help the visual controller to predict its position accurately and thereby avoid failure.

used techniques for anomaly detection.

1) *Prediction Error-based Labels*: We use a prediction error-based labeling scheme instead of the proposed HJ Reachability-based scheme to collect the training dataset for our AD. Specifically, if the TaxiNet prediction error is above a certain threshold for a particular image, then we label this image as an anomaly. We overlay the prediction error-based anomaly labels (blue) on top of the BRT-based labels (red) for one of the environments in Fig. 5(a).

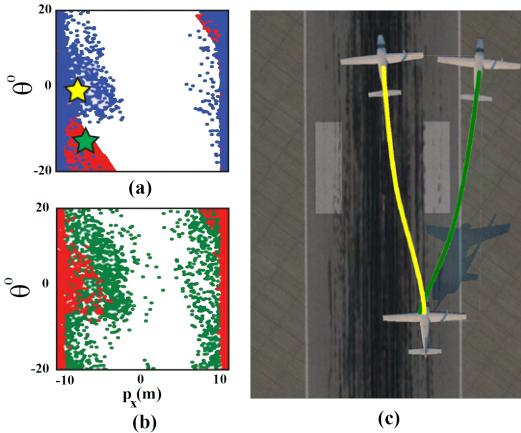


Fig. 5: (a) Comparison between prediction error (blue) and BRT-based (red) labels. (b) Prediction error-based labels for $threshold = 0.3$ (green) and $threshold = 0.6$ (red). (c) Yellow and Green lines show trajectories starting from the yellow and green stars, respectively.

It is evident that the prediction error-based labels may not be a good representative of the system-level failures. For example, the states near the Yellow star are anomalous as per prediction error but do not actually cause the system failure (the yellow trajectory in Fig. 5(c)), resulting in a pessimistic AD and hampering the system performance. On the other hand, certain states and images may have a small error from the TaxiNet module perspective (states near the Green star) and are not classified as anomalies; yet, they cascade to a system failure (the green trajectory in Fig. 5(c)). This results in an overly optimistic nature of the prediction of error-based AD in certain regions. This “unpredictable” nature of the prediction error-based AD persists with the change in the error threshold used for labeling (Fig. 5(b)). Unsurprisingly, we observe the same phenomenon in the AD trained on these prediction-error-based labels and, thus, omit those results for brevity. This leads us to conclude that component-level, prediction error-based labels may not be a good representative for determining system-level failures.

2) *Ensembled Predictions*: Another popular mechanism to detect anomalies is based on the predictive uncertainty of an ensemble of neural networks. To design an ensemble, we train 5 different versions of the TaxiNet with different weight initializations. If the variance between the predictions exceeds a threshold for any input image, we assign such an image as an anomaly. The corresponding anomalies over the statespace are shown in Fig. 6. We observe that this method does not perform well because the ensemble confidently makes incorrect predictions for some states (around the top right of the statespace), leading to faulty labels in those states. On the other hand, for some states (near the central region of the statespace), the ensemble disagreed on the predictions, leading to states being incorrectly marked as an anomaly. Hence, an ensembling-based approach also fails to accurately predict system-level anomalies.




Fig. 6: Labels generated using ensembling denoting failures (red) and success (white).

Fallback Mechanism. Equipped with a capable AD, we designed a simple fallback mechanism to ensure our aircraft’s safety under anomalous inputs. If at any point in its trajectory, the aircraft observes an image I that is classified as an anomaly (i.e., $\sigma(I) = 1$), the linear velocity of the plane (v in Eqn. (2)) is reduced by $0.01m/s$. Intuitively, this results in slowing the aircraft every time it encounters an anomaly, ultimately coming to a complete stop if it continues to encounter anomalies. As soon as the aircraft detects an image that is non-anomalous, the default TaxiNet controller takes over. Fig. 7 shows trajectories of the aircraft under the resultant controller (red trajectories). The nominal trajectories without the fallback controller are shown in dashed black. In the first case (Fig. 7 (a)), the anomaly is triggered due to the aircraft approaching too close to the runway boundary, while in the second case (Fig. 7 (b)) the anomaly is triggered due to semantic failure of the runway markings. In both cases, the TaxiNet controller leads the system off the runway. On the other hand, the fallback controller decreases the aircraft velocity whenever an anomaly is triggered to ensure system safety. This can also be seen from the velocity variation along the red trajectories. To further illustrate the advantage of the fallback controller, we compute the system BRT under the Taxinet controller and the fallback controller pipeline for one of the test cases on KMWH runway (Fig.7(c)). The BRT

(i.e., the set of failure states) is much smaller under the safety pipeline (the BRT volume decreases from 28.23% to 16.86%), showing that the proposed mechanism significantly reduces the number of closed-loop system failures. Finally,




Fig. 7: (a, b) Trajectory followed by the aircraft under the TaxiNet controller (dashed black line) and the safety pipeline (red line). The color shift in the red curve shows velocity variation due to the fallback controller. (c) The grey region represents the system BRT under the TaxiNet controller, and the blue region represents the BRT under the safety pipeline. The BRT obtained using the AD and the fallback controller is appreciably smaller than the one obtained using vanilla TaxiNet. (d) Input image at the start state in (a), causing system failure due to runway boundaries. (e) Input image at the start state in (b), causing system failure due to runway markings.

with the inclusion of the AD and the fallback controller along with the TaxiNet CNN, the system was able to run at 25Hz as compared to 33Hz when the TaxiNet governs the system as the sole controller. This allows us to run the system online without any severe latency issues.




Fig. 8: (a) Red dots represent incorrect predictions of the AD. The incorrect predictions are concentrated around the BRT boundary. (b) Input image corresponding to the green star in (a), incorrectly classified as failure by AD. (c) Input image corresponding to the yellow star in (a), incorrectly classified as success by AD. Test runway contains yellow-colored lines that confuse the TaxiNet, but such anomalies are not predicted well, as no runway with yellow lines is present in the training dataset.

Failure Modes of AD. Even with all these exciting discoveries, the AD fails to detect some anomalies. We find that this happens particularly for the images corresponding to the states near the BRT boundary. For example, in Fig. 8(b) and

8(c), we show two visually similar images in our test dataset (corresponding to the green and yellow stars in Fig. 8(a)). Even though the two images are visually similar, one image is inside the BRT and leads to the system failure (hence an anomaly), while the other does not. Such similar images with minor differences are hard to detect for our AD. Finally, we noticed that some of the semantics of the test environment that cause anomalies are not present in the training dataset and are not predicted well by the proposed AD (Fig. 8(d)). Such issues highlight the need for a continual update of anomaly detectors as more data about the system anomalies is obtained.

Incremental Training. A complementary approach to employing an AD can be to perform targeted incremental training of the TaxiNet model on the collected anomalous data. Essentially, these labeled anomalies represent scenarios where TaxiNet fails to perform optimally. Therefore, by conducting specialized training on these failure instances, we aim to fortify TaxiNet’s robustness in handling such cases. We present some preliminary result of incremental training of TaxiNet in Fig. 9. We plan on exploring this method further in future work.




Fig. 9: (Left) BRT slices for a DTP of 100m at 5 PM on KMWH runway showing the improvements due to incremental training (blue) over the vanilla TaxiNet (grey). The state depicted by the magenta star is only included in the grey region. (Right) Trajectory simulated from this state is seen to succeed under the incrementally trained TaxiNet (yellow) but fails under the vanilla version (green).

VI. DISCUSSION AND FUTURE WORK

In this work, we present an approach aimed at identifying and mitigating system-level anomalies of autonomous systems that rely on vision-based controllers for decision-making. By leveraging insights from reachability analysis, our approach learns an anomaly detector that effectively tackles concerns related to system-level safety during runtime. The learned anomaly detector is combined with a fallback controller to reduce the potential for catastrophic system failures significantly.

In the future, we plan to test our approach over a wide range of real-world scenarios and focus on more intricate fallback mechanisms that can efficiently work for different system-level failures. In addition, to further enhance the performance of the AD, one can take advantage of temporal information, such as the visual history, instead of only the current image. Finally, instead of relying on a classifier, other approaches could explore unsupervised methods like clustering via support vector machines or k-means applied to our annotated dataset.

REFERENCES

- [1] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.
- [2] Q. M. Rahman, P. Corke, and F. Dayoub, "Run-time monitoring of machine learning for robotic perception: A survey of emerging trends," *IEEE Access*, vol. 9, pp. 20 067–20 075, 2021.
- [3] R. Sinha, A. Sharma, S. Banerjee, T. Lew, R. Luo, S. M. Richards, Y. Sun, E. Schmerling, and M. Pavone, "A system-level view on out-of-distribution data in robotics," *arXiv preprint arXiv:2212.14020*, 2022.
- [4] L. Sun, X. Jia, and A. D. Dragan, "On complementing end-to-end human behavior predictors with planning," *arXiv preprint arXiv:2103.05661*, 2021.
- [5] A. Ben-Tal, D. den Hertog, A. D. Waegenaere, B. Melenberg, and G. Rennen, "Robust solutions of optimization problems affected by uncertain probabilities," *Management Science*, vol. 59, no. 2, pp. 341–357, 2013. [Online]. Available: <http://www.jstor.org/stable/23359484>
- [6] J. C. Duchi and H. Namkoong, "Learning models with uniform performance via distributionally robust optimization," *The Annals of Statistics*, vol. 49, no. 3, pp. 1378 – 1406, 2021. [Online]. Available: <https://doi.org/10.1214/20-AOS2004>
- [7] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2017, pp. 23–30.
- [8] M. Salehi, H. Mirzaei, D. Hendrycks, Y. Li, M. H. Rohban, and M. Sabokrou, "A unified survey on anomaly, novelty, open-set, and out-of-distribution detection: Solutions and future challenges," *arxiv preprint arxiv:2110.14051*, 2021. [Online]. Available: <https://arxiv.org/abs/2110.14051>
- [9] L. Ruff, J. R. Kauffmann, R. A. Vandermeulen, G. Montavon, W. Samek, M. Kloft, T. G. Dietterich, and K.-R. Müller, "A unifying review of deep and shallow anomaly detection," *Proceedings of the IEEE*, vol. 109, no. 5, pp. 756–795, 2021.
- [10] J. Yang, K. Zhou, Y. Li, and Z. Liu, "Generalized out-of-distribution detection: A survey," *arXiv preprint arXiv:2110.11334*, 2021.
- [11] V. Narayanan and R. B. Bobba, "Learning based anomaly detection for industrial arm applications," in *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy*, 2018, pp. 13–23.
- [12] T. Ji, A. N. Sivakumar, G. Chowdhary, and K. Driggs-Campbell, "Proactive anomaly detection for robot navigation with multi-sensor fusion," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4975–4982, 2022.
- [13] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi, "A review of uncertainty quantification in deep learning: Techniques, applications and challenges," *Information Fusion*, vol. 76, pp. 243–297, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1566253521001081>
- [14] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Curran Associates, Inc., 2017. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/file/9ef2ed4b7fd2c810847ffa5fa85bce38-Paper.pdf>
- [15] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *Proceedings of The 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. F. Balcan and K. Q. Weinberger, Eds., vol. 48. New York, New York, USA: PMLR, 20–22 Jun 2016, pp. 1050–1059. [Online]. Available: <https://proceedings.mlr.press/v48/gal16.html>
- [16] P. Zhang, J. Wang, A. Farhadi, M. Hebert, and D. Parikh, "Predicting failures of vision systems," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2014, pp. 3566–3573.
- [17] A. Sharma, N. Azizan, and M. Pavone, "Sketching curvature for efficient out-of-distribution detection for deep neural networks," in *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, ser. Proceedings of Machine Learning Research, C. de Campos and M. H. Maathuis, Eds., vol. 161. PMLR, 27–30 Jul 2021, pp. 1958–1967. [Online]. Available: <https://proceedings.mlr.press/v161/sharma21a.html>
- [18] A. Amini, W. Schwarting, A. Soleimany, and D. Rus, "Deep evidential regression," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 14 927–14 937. [Online]. Available: <https://proceedings.neurips.cc/paper/2020/file/aab085461de182608ee9f607f3f7d18f-Paper.pdf>
- [19] K. Chakraborty and S. Bansal, "Discovering closed-loop failures of vision-based controllers via reachability analysis," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2692–2699, 2023.
- [20] S. M. Katz, A. L. Corso, C. A. Strong, and M. J. Kochenderfer, "Verification of image-based neural network controllers using generative models," in *DASC*, 2021, pp. 1–10.
- [21] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-Jacobi reachability: A brief overview and recent advances," 2017.
- [22] I. Mitchell, A. Bayen, and C. J. Tomlin, "A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games," *TAC*, vol. 50, no. 7, pp. 947–957, 2005.
- [23] I. Mitchell and C. J. Tomlin, "Level set methods for computation in hybrid systems," in *HSCC*. Springer, 2002, pp. 310–323.
- [24] I. M. Mitchell *et al.*, "A toolbox of level set methods," *UBC Department of Computer Science Technical Report TR-2007-11*, p. 31, 2007.
- [25] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International conference on machine learning*. PMLR, 2019, pp. 6105–6114.