

Testing Rare Downstream Safety Violations via Upstream Adaptive Sampling of Perception Error Models

Craig Innes and Subramanian Ramamoorthy^{1,2}

Abstract—Testing black-box perceptual-control systems in simulation faces two difficulties. Firstly, perceptual inputs in simulation lack the fidelity of real-world sensor inputs. Secondly, for a reasonably accurate perception system, encountering a rare failure trajectory may require running infeasibly many simulations. This paper combines perception error models—surrogates for a sensor-based detection system—with state-dependent adaptive importance sampling. This allows us to efficiently assess the rare failure probabilities for real-world perceptual control systems within simulation. Our experiments with an autonomous braking system equipped with an RGB obstacle-detector show that our method can calculate accurate failure probabilities with an inexpensive number of simulations. Further, we show how choice of safety metric can influence the process of learning proposal distributions capable of reliably sampling high-probability failures.

I. INTRODUCTION

Perceptual-control systems, such as in automated vehicles (AV), must be rigorously evaluated before deployment. Safety specifications can be defined formally in Signal Temporal Logic (STL) [1], [2], but the black-box nature and operating scenario complexity of such systems often preclude analytical guarantees. Instead, we seek assurances via testing.

While one can aim at the ideal of conducting exhaustive real world tests, this is infeasible [3]. In addition to the cost, these tests represent danger to the public. Alternatively, we can simulate our system, but simulation poses two core problems: First, while modern simulators can replicate real-world physics with high accuracy, perceptual inputs from a simulated sensor are far from their real-world counterparts. Figure 1 shows the photorealism gap between a real and simulated RGB-camera image. In a black-box system with a mixture of perception and control components, this means that, if we train the upstream perception components on simulated inputs, we lack assurances that simulated downstream control failures will correspond to failures experienced by our system when using real-world sensor inputs.

However, a key insight is that, to trust our simulation, photorealism may be unnecessary. It is sufficient that the simulated perception system is accurate in the same situations as the real one. *Perception Error Models* (PEMs) [4] are surrogate models of a perception system which, instead




Fig. 1: The difference in fidelity between real traffic images (top) and simulated images rendered in CARLA (bottom).

of taking sensor inputs, take salient characteristics of those inputs, outputting the expected accuracy of the target system.

A second problem is, for any modestly accurate perception system, safety violations can be extremely rare – especially if a system must fail many times in sequence to cause danger. Consider an automated braking system with an image-based obstacle detector. If the detector achieves 99% accuracy, and needs to miss an obstacle for one second to cause a crash, then at 20 images-per-second, failure probability is $(1.0 - 0.99)^{20}$. This means we would expect to run 1×10^{40} simulations before observing a *single crash*.

Instead of direct sampling, *importance sampling* draws from a proposal distribution where failures are likely, then re-weights samples according to their likelihood under the target. The challenge is finding proposals where failures are likely, while producing trajectories common under the original distribution. Existing importance sampling approaches define proposal distributions over full trajectories, often failing to produce proposals capable of generating extremely rare events caused by long sequences of perception failures.

This paper proposes a method for estimating black-box perceptual-control system failures against STL specifications in simulation. To achieve this, we first train a PEM to mimic a perception system trained on real-world images. Then, we exploit this PEM to drive learning of a low-variance sampler using *state-dependent* adaptive importance sampling.

We apply our method to an automated braking scenario in the CARLA simulator. The perceptual-control system under test uses a YOLO obstacle detector [5] trained on images from the KITTI dataset [6]. Our experiments show that, with a representative set of training images and sufficiently descriptive salient variables, our method accurately estimates how upstream perception accuracy affects downstream safety violations. Since multiple robustness metric choices exist for STL, our experiments also demonstrate how different metrics

¹Craig Innes (craig.innes@ed.ac.uk) and Subramanian Ramamoorthy (s.ramamoorthy@ed.ac.uk) are with the School of Informatics, University of Edinburgh, Edinburgh, United Kingdom.

² Work supported by a grant from the UKRI Strategic Priorities Fund to the UKRI Research Node on Trustworthy Autonomous Systems Governance and Regulation (EP/V026607/1, 2020-2024). For the purpose of open access, the author(s) has applied a Creative Commons Attribution (CC BY) license to any Accepted Manuscript version arising

affect the learning process of adaptive sampling.

II. ESTIMATING THE PROBABILITY OF FAILURE

We first formally describe failure probability estimation for our perceptual-control system: Our environment contains states s in d -dimensional space $\mathcal{S} = \{(x_1, \dots, x_d) \mid x_{1\dots d} \in X_{1\dots d}\}$. Our perceptual-control system $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ will, given state s , output action $a \in \mathcal{A}$ with probability $\pi(a|s)$. In the AV domain, variables X_1, \dots, X_d may include vehicle position/velocity, or pixel values from a camera; action space \mathcal{A} may consist of whether an obstacle detection system detects a vehicle in frame.

The probability of transitioning to state s' from s is $p(s', a \mid s) = \pi(a \mid s)p(s' \mid s, a)$. This paper makes the simplifying assumption (as in [7]) that the only source of stochasticity is perception system outputs. E.g., once the perception system decides an obstacle is ahead, the decision to brake (and effect of braking on velocity) is deterministic:

$$p(s', a \mid s) = \pi(a \mid s) \quad (1)$$

Starting from s_0 we define a T -step *simulation rollout* as $\tau = [s_0, a_1, s_1, a_2 \dots s_{T-1}]$, where the rollout probability is:

$$p(\tau) = \prod_{t=1}^{T-1} \pi(a_t \mid s_{t-1}) \quad (2)$$

We define a *safety metric* $r(\varphi, \tau) \in \mathbb{R}$, where φ is a specification, and $r(\varphi, \tau) < 0$, denotes violation. Our goal is to compute the probability our system falls below some safety threshold γ :

$$\mathbb{E}_\pi [\mathbb{1} \{r(\tau, \varphi) \leq \gamma\}] = \int_\tau \mathbb{1} \{r(\tau, \varphi) \leq \gamma\} p(\tau) \quad (3)$$

Subsequent sections describe the components needed to estimate (3) — specifying φ and $r(\tau, \varphi)$, constructing a PEM, and learning a proposal distribution.

A. Specifying Safety Requirements with STL

Signal Temporal Logic (STL) is a language for expressing properties of timed, continuous signals [8]. Its grammar is:

$$\varphi := \top \mid \eta \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \mid \square_I \varphi \mid \diamond_I \varphi \mid \varphi_1 \mathcal{U}_I \varphi_2 \quad (4)$$

where η is any predicate $\rho(s) - b \leq 0$ (with $b \in \mathbb{R}$ and $\rho : \mathcal{S} \rightarrow \mathbb{R}$). Symbol $\square_I \varphi$ means φ is *always* true at every time-step within interval I . Symbol $\diamond_I \varphi$ means φ must *eventually* be true at some time-step in I , and $\varphi_1 \mathcal{U}_I \varphi_2$ means φ_1 must remain true within I *until* φ_2 becomes true. For example:

$$\square_{[0, T]} (\|C_{ego}.pos - C_{other}.pos\| \geq 2.0) \quad (5)$$

means “From 0 to T , the distance between the ego car and other car must never drop below 2 metres”. STL formulae are paired with a *robustness metric* $r(\tau, \varphi) \in \mathbb{R}$, specifying how strongly τ satisfied φ ($r(\tau, \varphi) > 0$) or violated φ ($r(\tau, \varphi) < 0$). Classical robustness focuses on *spatial*

robustness, applying hard minimums across conjunctions [8]. For instance, equation (6) evaluates the *always* operator:

$$r(\tau, \square_{[0, T]} \varphi) = \min(r(s_0, \varphi), \dots, r(s_T, \varphi)) \quad (6)$$

Applied to (5), this corresponds to the single state in τ with the smallest distance between vehicles.

Alternative robustness metrics exist, each offering trade-offs in time/space robustness, soundness, and smoothness: AGM-robustness [9] replaces the min/max functions in *always/eventually* with geometric/arithmetic means; Smooth Cumulative Robustness [10], combines smooth-approximations to min/max with cumulative summations.

B. Perception Error Models

In equation (2), τ is determined by our perceptual model $\pi(\mathcal{A} \mid s)$. For an obstacle detector, this could be a function f over RGB images, outputting *detected* or *-detected*:

$$\pi(\text{detected} \mid s) = f_\theta(s) \quad (7)$$

Here θ could represent e.g., the weights of a Convolutional Neural Network trained on real traffic images. When testing in simulation however, we cannot use f_θ directly, since our simulator does not produce real sensor inputs. We could train an alternative f based on rendered images, but as discussed in the introduction, such images correspond poorly to the real world. The added cost of including a rendering and image-processing step in the simulation pipeline also increases simulation costs by an order of magnitude.

Here is the key insight of PEMs: For simulated risk assessment to be useful, it is sufficient that, for the same state s , our surrogate \hat{f} and function f share the same error:

$$\forall s \in \mathcal{S}, f_\theta(s) = \hat{f}_\theta(g(s)) \quad (8)$$

where g is a projection of \mathcal{S} (which could contain RGB image pixels, or other sensor data), down to a set of *salient variables* — a subset of perceptual features necessary to predict system performance. This may include obstacle distance to camera, vehicle type, or level of occlusion.

We train \hat{f} using m training samples D , consisting of predictions from our perception system f paired with the corresponding salient variables:

$$D = \{\langle f(s_0), g(s_0) \rangle, \dots, \langle f(s_m), g(s_m) \rangle\} \quad (9)$$

We then find the $\hat{\theta}$ which minimizes the binary cross entropy (BCE) between predictions $f_\theta(s)$ and $\hat{f}_\theta(g(s))$:

$$\sum_{i=0}^m f(s_i) \log [\hat{f}_{\hat{\theta}}(g(s_i))] + (1 - f(s_i)) \log [1 - \hat{f}_{\hat{\theta}}(g(s_i))] \quad (10)$$

Once \hat{f}_θ is learned, we can use it as a drop-in surrogate for f in our simulator by projecting each s down to its salient variables as they arrive online.

C. Adaptive Importance Sampling

Given N simulations, we can estimate (3) via monte-carlo approximation:

$$\hat{\mu}_{\text{MC}} = \frac{1}{N} \sum_{i=0}^N \mathbb{1}_{\{r(\tau_i, \varphi) \leq \gamma\}} \quad (11)$$

For rare-events, this is highly inefficient. For N samples, the relative error [11] is approximately $\sqrt{(N\mu)^{-1}}$. This means to achieve a 1% relative error with event probability $\mu = 10^{-6}$, we need to run 10^{10} simulations. A more efficient approach is to use *importance sampling* to draw from a proposal q where the rare-event is more likely, then re-weight each sample with respect to target distribution p :

$$\hat{\mu}_{\text{IS}} = \frac{1}{N} \sum_{i=0}^N \mathbb{1}_{\{r(\tau_i, \varphi) \leq \gamma\}} \frac{p(\tau_i)}{q(\tau_i)} \quad (12)$$

Rather than defining distributions over τ , our modified importance sampler defines proposals over *states* $s \in \mathcal{S}$:

$$\frac{1}{N} \sum_{i=0}^N \mathbb{1}_{\{r(\tau_i, \varphi) \leq \gamma\}} \prod_{t=0}^T \frac{\hat{f}_{\hat{\theta}}(g(s_{i,t}))}{q_{\phi}(h(s_{i,t}))} \quad (13)$$

We use $s_{i,t}$ to denote state at time t of the i -th simulation. Following (2) and (7), we factor $p(\tau)$ as a product of state-dependent outputs from PEM \hat{f} . Likewise, $q(\tau)$ is factored as product of probabilities $q_{\phi}(h(s))$. Note, $h(s)$ does not need to be the same as $g(s)$. $h(s)$ might project s over a small subset of features we suspect influence safety e.g., vehicle acceleration or velocity.

We want a proposal where failure events are common, and trajectories have high likelihood under the target distribution. The perfect proposal q^* has the form:

$$q^*(\tau) = \left(\frac{p(\tau)}{\mu} \right) \mathbb{1}_{r(\tau, \varphi) < \gamma} \quad (14)$$

Of course, we do not know μ , but given example trajectories $[\tau_0, \dots, \tau_N]$, we can learn good proposal parameters ϕ by minimizing the KL-divergence between the data and (14):

$$- \sum_{i=0}^N w_i \left(\sum_{t=0}^T \log q_{\phi}(h(s_{i,t})) \right) \mathbb{1}_{r(\tau_i, \varphi) < \gamma} \quad (15)$$

with weights w_i defined as:

$$w_i = \prod_{t=0}^T \frac{\hat{f}_{\hat{\theta}}(g(s_{i,t}))}{q_{\phi}(h(s_{i,t}))} \quad (16)$$

However, we have a problem: If no trajectory in τ_0, \dots, τ_N passes threshold $\mathbb{1}_{r(\tau, \varphi) < \gamma}$, this halts optimization. To address this, we apply techniques from the cross-entropy method (CEM) literature [11] — Instead of filtering all simulations by threshold γ , we iteratively sample multiple batches of simulations. At each stage $\kappa \in [0, \dots, K]$, we run N_{κ} simulations, then sort $\{\tau_1, \dots, \tau_{N_{\kappa}}\}$ by $r(\tau, \varphi)$. In place of γ , we set an intermediate threshold $\gamma_{\kappa} = \max(\gamma, r(\tau_{\sigma N_{\kappa}}))$, where σ is a desired quantile within range $0.95 \leq \sigma < 1.0$.

This ensures some simulations meet the threshold at every stage, driving incremental proposal learning via (15).

Because we are finding a proposal via numerical optimization of a product of state dependent distributions (rather than e.g., a single exponential distribution admitting a closed-form solution [12], [13], [14]), and because low safety thresholds are characterized by long chains of consecutive failures, we have another problem: The difference between the highest and lowest trajectory weights w may become enormous. This difference can lead to numerical instability, with less likely trajectories having effectively zero weight during learning. If those less likely trajectories are necessary to achieve lower safety thresholds, this can result in the optimization stalling at an intermediate threshold γ_{κ} .

Defensive Importance Sampling techniques can be used to limit the ratio between the target and proposal likelihoods for a single state [15], but such techniques place no limit on weights w_i resulting from the *product* of likelihood sequences. Instead, our approach is to apply a smoothing parameter $\alpha \in [0, 1]$ to importance weights:

$$- \sum_{i=0}^N w_i^{\alpha} \left(\sum_{t=0}^T \log q_{\phi}(h(s_{i,t})) \right) \mathbb{1}_{r(\tau_i, \varphi) < \gamma} \quad (17)$$

This lessens the gap between the highest and lowest weights, allowing our adaptive learning process to progress.

Algorithm 1 Simulated Risk Estimation with PEMs

- 1: **function** ADAPTIVE-EST($D, \varphi, N_{\kappa}, N_e, \gamma$)
 - 2: $\hat{\theta} \leftarrow$ Train PEM using D with (10)
 - 3: $\phi \leftarrow$ Initialize proposal distribution
 - 4: **for** $\kappa = 1$ to K **do**
 - 5: $\{\tau_0 \dots \tau_{N_{\kappa}}\} \leftarrow$ Sample N_{κ} rollouts with q_{ϕ}
 - 6: Sort $\{\tau_0 \dots \tau_{N_{\kappa}}\}$ by $r(\tau, \varphi)$
 - 7: $\gamma_{\kappa} \leftarrow \max(\gamma, r(\tau_{\lfloor \sigma N_{\kappa} \rfloor}))$
 - 8: $\phi \leftarrow$ Min (17) with $\gamma_{\kappa}, q_{\phi}, \hat{f}_{\hat{\theta}}, \{\tau_0 \dots \tau_{N_{\kappa}}\}$
 - 9: $\{\tau_0 \dots \tau_{N_e}\} \leftarrow$ Sample N_e rollouts with q_{ϕ}
 - 10: **return** $\hat{\mu}_{\text{IS}}$ using $\{\tau_0 \dots \tau_{N_e}\}$ with (13)
-

Algorithm (1) Combines our STL robustness components, PEM, and adaptive sampling to give our full method. After training our PEM to replicate real-world perceptual system behaviour as outlined in section (II-B), we perform K stages of black-box simulation. At each stage κ , we conduct N_{κ} simulations, set threshold γ_{κ} based on the upper quantile of rollouts, then use these rollouts to optimize parameters ϕ for the next iteration. After the K iterations, we perform a final N_e simulations to estimate failure probability μ .

III. EXPERIMENTS

Our experiments apply algorithm (1) to an automated braking problem in the CARLA vehicle simulator [16]. We show our method can estimate the probability of a rare collision event efficiently, yielding a greater number of high-probability failure events (compared with classical monte-carlo sampling, or a naive custom proposal). Further, we



Fig. 2: Automated braking scenario — When the front car abruptly stops, our vehicle must brake to avoid crashing.

show how the choice of STL robustness metric can affect learning of the proposal distribution.

Figure (2) illustrates our scenario: Our car follows another for $T = 100$ time steps. At the end of the road, the front vehicle brakes, and our car must brake to avoid crashing. We use (5) as our safety condition — If the distance between vehicles is ever less than 2 metres we consider it a “crash”¹.

Our vehicle’s perceptual-control system comprises two main components: a physical controller and a perception system. The physical controller is a simplification of the `BasicAgent` from the base CARLA library²: A global planner sets road waypoints for the car to follow; A local planner applies PID control to reach max speed if the road is obstacle-free, and applies emergency brakes if there is an obstacle within 10 metres of the vehicle³. The perception system is a neural-network-based obstacle detector trained on RGB images of real traffic from the KITTI dataset [6]. We first describe how we train a PEM from this detector, then outline its use within CARLA.

A. Learning a PEM of the KITTI dataset

KITTI is an object detection dataset of 1242×375 traffic images, with labels denoting vehicle type and 2D bounding boxes. It contains 7481 training and 7518 test images.

The obstacle detection system uses YOLO-v3 [5], trained on KITTI. Figure 3 sketches YOLO’s operation. Images are divided into regions, bounding boxes produced, and detection probabilities $f(s)$ predicted in one pass through the network.

We next learn a PEM to act as a surrogate \hat{f}_{θ} for our YOLO perception system. Figure (4) shows our feed-forward PEM architecture. To train, we first pre-process the data from our YOLO-KITTI detector, separating individual obstacles and pairing them to their closest bounding box prediction via Hungarian Matching [17]. For salient variables $g(s)$, we use 14 dimensions based on auxiliary labels provided by the KITTI dataset for each image. The 14 dimensions are: A 6-dimensional one-hot encoding of vehicle category, a 3-dimensional one-hot encoding for the amount the vehicle is occluded (None, Partial, Mostly), the xyz obstacle location, and the rotation of the obstacle on the y -axis. After removing obstacles with unknown values, we were left with 37498

¹Note, the true distance is less than 2 metres, as we measure from the windscreen camera to the rear of the other car [6]

²<https://carla.readthedocs.io>

³Code at <https://github.com/craigiedon/CarlaStuff>

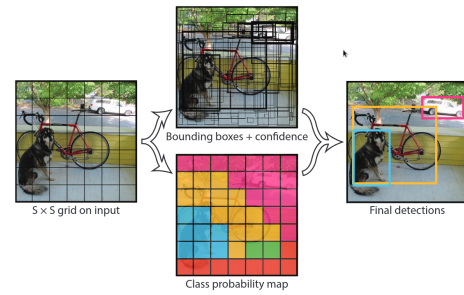


Fig. 3: (Diagram from [5]) The YOLO detector. Images are divided into a grid. For each grid cell, it predicts B bounding boxes, box confidences, and class probabilities.




Fig. 4: PEM architecture of 3 feed-forward layers (20-hidden units each), ReLU activations and Batch Normalization.

obstacles to train over. We then minimize loss function (10) using the Adam Optimizer [18] over 300 epochs.

Note, our PEM training has two limits. First, by considering labels on single obstacles, our model is unable to capture false positive detections. This is a common limitation [4], but could be overcome with a more expressive surrogate. Second, we assume the KITTI auxiliary labels are expressive enough to explain variance in detection performance across the domain. This is likely untrue—perception systems are affected by additional factors such as time of day, weather conditions, road surface reflections etc. In an industry-scale system deployment, we would expect to have dedicated teams of annotators to design and record a more widely expressive set of salient variables for each image. These issues lie outside the scope of this paper.

Table (I) summarizes our PEM’s performance (ML-NN), via 5-fold cross-validation. Several baselines are also given: LR fits a logistic regressor; $\text{Guess-}\mu$ returns the mean detection probability. We also fit a complex Bayesian Neural Network (B-NN). Rather than finding the most likely $\hat{\theta}$, B-NN learns a distribution over parameters using variational inference [19]. We report the binary cross-entropy (BCE) and area under the receiver operating characteristic curve (ROC-AUC) for each model. Across baselines, our original model ML-NN performs best.

It is perhaps surprising that B-NN does not perform better. It may be that the variational family used to capture the detection distribution was not sufficiently expressive. Alternatively, held-out data used in the 5-fold validation may be insufficiently different from the training data to realise the benefit of marginalizing over parameters. While not the focus of this paper, it could be interesting to investigate the performance of such PEMs in cases of distribution shift and out-of-domain images.

TABLE I: Results of 5-fold cross validation across perception error models for KITTI obstacle detection task.

Model	BCE	ROC-AUC
B-NN	0.238	0.726
LR	0.244	0.688
ML-NN	0.226	0.777
Guess- μ	0.260	0.5

TABLE II: Estimated failure probabilities under different sampling strategies. Averaged over 10 repetitions, with standard errors given in brackets.

Method	$\hat{\mu}$	Failures	NLL
MC ₁₀₀	0.0	0 / 100	-
MC ₁₀₀₀₀	0.0	0 / 10000	-
NAIVE-50 ₁₀₀	2.64e-48 (\pm 4.69e-49)	77 / 100	199.95
NAIVE-50 ₁₀₀₀	4.75e-44 (\pm 2.16e-44)	748 / 1000	177.97
ADAPTIVE-PEM ₁₀₀	5.02e-15 (\pm 3.83e-15)	52/100	41.6

B. Adaptive Sampling of Rare Crash Events

We now use our PEM to drive adaptive sampling in algorithm (1). We run $K = 10$ stages, with $N_\kappa = 100$, and $N_e = 100$. Each simulation runs for $T = 100$ time steps, with a delta 0.05. Cars start at 70 km/h, 15 metres apart.

Proposal network q_ϕ is a two-layer feed-forward network, with 32 hidden nodes per layer, and *ReLU* activations. Transformation function $h(s)$ projects s down to a single dimension—the distance (in metres) between vehicles. To ensure our initial proposal produces trajectories similar to the original PEM, we pre-train ϕ with 100 simulations from \hat{f}_θ . In adaptation stages, we minimize (17) with $\alpha = 0.1$. We use the Adam optimizer [18] for 500 epochs.

We compare our method (ADAPTIVE-PEM₁₀₀) to several baselines: MC₁₀₀ and MC_{10,000} estimate crash probability using 100 and 10000 raw monte-carlo samples respectively. We also compare against a naive custom importance sampler (NAIVE-50₁₀₀ and NAIVE-50₁₀₀₀) which ignores state information, instead detecting obstacles with a flat 50% probability. This is likely to produce more frequent crashes, but the trajectories may be unlikely under the target.

For our robustness metrics $r(\tau, \varphi)$ used to evaluate thresholds $\gamma\kappa$, we consider three variants—Classical Spatial Robustness (r_c), AGM Robustness (r_a), and Smooth-Cumulative Robustness (r_s), with $\gamma = 0$ across all three.

C. Results

Table II shows the results of each sampling strategy. Clearly, raw monte-carlo sampling is insufficient for this problem — After 10,000 simulations, MC₁₀₀₀₀ has produced 0 failures, reporting 0% crash probability. This is unsurprising given the theoretical results around relative-error combined with the high accuracy of our target detector ($> 90\%$ average detection accuracy). Since conducting enough monte-carlo simulations to extract a ground-truth probability is computationally infeasible, we instead rely on two proxies—Proportion of final trajectories resulting in a crash, and average negative log likelihood (NLL) of failures.

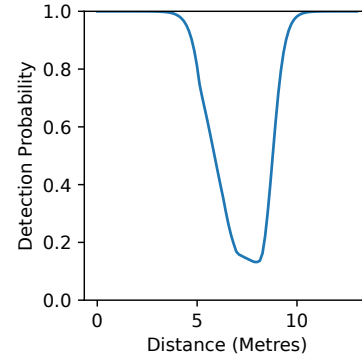


Fig. 5: Learned proposal at $\kappa = 10$ using Alg (1)

TABLE III: STL metrics comparison

Metric	$\hat{\mu}_{is}$	Failures	NLL
r_c	5.02e-15 (\pm 3.83e-15)	52/100	41.6
r_a	6.76e-15 (\pm 3.64e-15)	41/100	42.4
r_s	2.01e-15 (\pm 5.38e-16)	40/100	41.3

ADAPTIVE-PEM₁₀₀, successfully learns a proposal that produces more than half of its simulations as failures, with an average NLL of 41.6. In contrast, while NAIVE-50 produces a larger failure proportion, the NLL is far higher (199.95). This means trajectories from NAIVE-50 are 1×10^{69} times less likely under \hat{f}_θ than those from ADAPTIVE-PEM. As a result, NAIVE-50 is dominated by a few trajectories with large importance weights, leading to inaccurate estimates of $\hat{\mu}$. In fact, we can see that as the number of simulations increases from 100 to 1000, NAIVE-50’s estimate gradually moves towards ADAPTIVE-PEM’s estimate.

Figure (5) shows the detection rate of ADAPTIVE-PEM as a function of vehicle distance. At distances where safety is not a concern, the proposal mimics the detection accuracy of the original PEM (between 98–99.9%). As the cars approach distances where a failure to brake would cause a crash, the proposal gradually lowers the detection probability.




Fig. 6: Proposal distributions learned for r_c , r_a , and r_s at intermediate adaptation stage $\kappa = 3$. Averaged over 10 repetitions.




Fig. 7: Trajectories for the three lowest ranked simulation runs according to the r_c , r_a , and r_s robustness metrics.

For robustness metrics $r(\tau, \varphi)$, Table (III) shows that the final proposals and $\hat{\mu}$ -estimates after $K = 10$ adaptation stages is roughly similar across all three. However, as figure (6) reveals, the behaviour of the three metrics at intermediate stages can vary.

To understand why this occurs, consider the different ways r_c , r_a , and r_s rank the safety of a batch of simulations. Figure (7) shows the trajectories that each metric chose as “least safe” from a collection of 10000 random simulations. These choices align with our understanding of how each metric calculates robustness: r_c computes robustness as a strict minimum of distances across all time steps, so chooses the trajectory with the single smallest distance value. In contrast, r_a applies a geometric mean over distances, leading the metric to choose a trajectory with a (marginally) higher minimum distance, but which stays lower for a longer time-span. Finally, r_s applies a smoothed approximation to the minimum, leading it to choose a trajectory somewhere between the absolute minimum and mean values. These ranking differences affect trajectory sorting in line 6 of Algorithm (1), which then influences the trajectories chosen in each adaptation stage to optimize (17)

The main insight is, given limited adaptation budget, the choice of $r(\tau, \varphi)$ can influence threshold levels γ_κ , trajectory ordering, and potentially learning process behaviour.

IV. RELATED WORK

Many works test systems via black-box falsification or adaptive-stress testing [20], [21], [22], [23], [24], [25]. Such works focus on finding a single failure trajectory, rather than estimating the overall probability of failure.

This paper estimated rare-event failures by augmenting CEM techniques [11]. Other works use CEM for AV risk estimation [26], [14], [12]. These typically learn proposals over trajectories, an approach which can fail when rare events result from a product of probabilities over long-running trajectories. Additionally, these works often limit themselves to exponential-family proposals over a set of initial conditions; They avoid sampling perceptual behaviours of the vehicle itself. Our method instead learns proposals over individual states, using PEMS to handle sensor observations.

One state-dependent approach close to ours frames rare-event sampling as an approximate dynamic programming problem [7]. It requires the user to have prior knowledge of the transition model, and the ability to jump to arbitrary states. For black-box simulators like CARLA, neither requirement holds.

An alternative adaptive sampling technique is multi-level splitting [27], [28]. Here initial simulation trajectories are treated as particles which are perturbed using Markov-Chain Monte Carlo (MCMC). This method scales better to long-running sequential problems, but conducting MCMC transitions per trajectory per stage is intractable for simulators with non-trivial run time.

Continuation approaches use earlier perception system training epochs as a proxy for proposal stages [29]. They assume the user has access to earlier training checkpoints, and that failures in those earlier stages are similar to those encountered by the final system.

PEMS are applied in AV as surrogates for image, radar, and lidar sensors [4], [30], [31], [32]. The primary motivation of such applications is usually to bypass expensive rendering, providing fast, low-fidelity predictions. In contrast, our work takes advantage of the probabilistic outputs provided by PEMS to also drive adaptive importance sampling.

We make the core assumption that our PEM is *well calibrated* [33]. Several issues can challenge this assumption. First, our model may encounter *epistemic uncertainty* if images in simulation are far from the training distribution [34]. Second, our model may encounter *aleatoric uncertainty* if the salient variables insufficiently express relevant safety aspects of a scene [35]. In future work, we aim to explore this relationship between epistemic/aleatoric uncertainty in the upstream perceptual system, and the guarantees given on downstream failures.

There have been multiple attempts to encode complex traffic rules using temporal logic [36], [37], [1], [2]. In future work, we aim to scale our method from automated braking to more complex traffic scenarios, investigating the effect of balancing multiple competing rules.

For testing of black-box systems, one issue is assessing trust in your final estimation. Safety certificates help ensure risk is not under or over-estimated [38], [39], [40], [41]. These techniques are complementary to our task, and can be applied as an additional step after sampling.

V. CONCLUSION

This paper proposed a method for testing black box perceptual-control systems in simulation, even when the perception system takes real-world sensor inputs, and failures are rare. We used PEMS to drive state-dependent adaptive importance samplers, with an STL robustness metric to assess failure thresholds. Our experiments in automated braking showed our method efficiently estimates rare failure probabilities, learning distributions that generate a high percentage of high-likelihood failure trajectories.

REFERENCES

- [1] M. Hekmatnejad, S. Yaghoubi, A. Dokhanchi, H. B. Amor, A. Shrivastava, L. Karam, and G. Fainekos, "Encoding and monitoring responsibility sensitive safety rules for automated vehicles in signal temporal logic," in *Proceedings of the 17th ACM-IEEE International Conference on Formal Methods and Models for System Design*, 2019, pp. 1–11.
- [2] N. Arechiga, "Specifying safety of autonomous vehicles in signal temporal logic," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 58–63.
- [3] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 182–193, 2016.
- [4] J. Sadeghi, B. Rogers, J. Gunn, T. Saunders, S. Samangooei, P. K. Dokania, and J. Redford, "A step towards efficient evaluation of complex perception tasks in simulation," *arXiv preprint arXiv:2110.02739*, 2021.
- [5] J. Redmon and A. Farhadi, "Yolov3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.
- [6] A. Geiger, P. Lenz, C. Stillér, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [7] A. Corso, R. Lee, and M. J. Kochenderfer, "Scalable autonomous vehicle safety validation through dynamic programming and scene decomposition," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–6.
- [8] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer, 2010, pp. 92–106.
- [9] N. Mehdipour, C.-I. Vasile, and C. Belta, "Arithmetic-geometric mean robustness for control from signal temporal logic specifications," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 1690–1695.
- [10] I. Haghghi, N. Mehdipour, E. Bartocci, and C. Belta, "Control from signal temporal logic specifications with smooth cumulative quantitative semantics," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 4361–4366.
- [11] D. P. Kroese, R. Y. Rubinstein, and P. W. Glynn, "The cross-entropy method for estimation," in *Handbook of statistics*. Elsevier, 2013, vol. 31, pp. 19–34.
- [12] Y. Kim and M. J. Kochenderfer, "Improving aircraft collision risk estimation using the cross-entropy method," *Journal of Air Transportation*, vol. 24, no. 2, pp. 55–62, 2016.
- [13] D. Zhao, H. Lam, H. Peng, S. Bao, D. J. LeBlanc, K. Nobukawa, and C. S. Pan, "Accelerated evaluation of automated vehicles safety in lane-change scenarios based on importance sampling techniques," *IEEE transactions on intelligent transportation systems*, vol. 18, no. 3, pp. 595–607, 2016.
- [14] M. O' Kelly, A. Sinha, H. Namkoong, R. Tedrake, and J. C. Duchi, "Scalable end-to-end autonomous vehicle testing via rare-event simulation," in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [15] M. F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Miguez, and P. M. Djuric, "Adaptive importance sampling: The past, the present, and the future," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 60–79, 2017.
- [16] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *Conference on robot learning*. PMLR, 2017, pp. 1–16.
- [17] D. A. Forsyth and J. Ponce, *Computer vision: a modern approach*. prentice hall professional technical reference, 2002.
- [18] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.
- [19] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, "Hands-on bayesian neural networks—a tutorial for deep learning users," *IEEE Computational Intelligence Magazine*, vol. 17, no. 2, pp. 29–48, 2022.
- [20] T. Dreossi, S. Jha, and S. A. Seshia, "Semantic adversarial deep learning," in *International Conference on Computer Aided Verification*. Springer, 2018, pp. 3–26.
- [21] T. Dreossi, A. Donzé, and S. A. Seshia, "Compositional falsification of cyber-physical systems with machine learning components," *Journal of Automated Reasoning*, vol. 63, no. 4, pp. 1031–1053, 2019.
- [22] Y. Annpureddy, C. Liu, G. Fainekos, and S. Sankaranarayanan, "S-taliro: A tool for temporal logic falsification for hybrid systems," in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2011, pp. 254–257.
- [23] K. D. Julian, R. Lee, and M. J. Kochenderfer, "Validation of image-based neural network controllers through adaptive stress testing," in *2020 IEEE 23rd international conference on intelligent transportation systems (ITSC)*. IEEE, 2020, pp. 1–7.
- [24] R. Lee, O. J. Mengshoel, A. Saksena, R. W. Gardner, D. Genin, J. Silbermann, M. Owen, and M. J. Kochenderfer, "Adaptive stress testing: Finding likely failure events with reinforcement learning," *Journal of Artificial Intelligence Research*, vol. 69, pp. 1165–1201, 2020.
- [25] H. Delecki, M. Itkina, B. Lange, R. Senanayake, and M. J. Kochenderfer, "How do we fail? stress testing perception in autonomous vehicles," *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5139–5146, 2022.
- [26] D. Zhao, X. Huang, H. Peng, H. Lam, and D. J. LeBlanc, "Accelerated evaluation of automated vehicles in car-following maneuvers," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 733–744, 2017.
- [27] J. Norden, M. O'Kelly, and A. Sinha, "Efficient black-box assessment of autonomous vehicle safety," *arXiv preprint arXiv:1912.03618*, 2019.
- [28] A. Sinha, M. O'Kelly, R. Tedrake, and J. C. Duchi, "Neural bridge sampling for evaluating safety-critical autonomous systems," *Advances in Neural Information Processing Systems*, vol. 33, pp. 6402–6416, 2020.
- [29] J. Uesato, A. Kumar, C. Szepesvari, T. Erez, A. Ruderman, K. Anderson, K. Dvijotham, N. Heess, and P. Kohli, "Rigorous agent evaluation: An adversarial approach to uncover catastrophic failures," in *International Conference on Learning Representations*, 2019.
- [30] P. Mitra, A. Choudhury, V. R. Aparow, G. Kulandaivelu, and J. Dauwels, "Towards modeling of perception errors in autonomous vehicles," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 3024–3029.
- [31] A. Piazzoni, J. Cherian, M. Slavik, and J. Dauwels, "Modeling perception errors towards robust decision making in autonomous vehicles," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 2021.
- [32] V. Berkahn, M. Kleiber, J. Langner, C. Timmermann, and S. Weber, "Traffic dynamics at intersections subject to random misperception," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 5, pp. 4501–4511, 2021.
- [33] J. Gawlikowski, C. R. N. Tassi, M. Ali, J. Lee, M. Humt, J. Feng, A. Kruse, R. Triebel, P. Jung, R. Roscher, et al., "A survey of uncertainty in deep neural networks," *arXiv preprint arXiv:2107.03342*, 2021.
- [34] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," *Advances in neural information processing systems*, vol. 31, 2018.
- [35] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" *Advances in neural information processing systems*, vol. 30, 2017.
- [36] L. Gressenbuch and M. Althoff, "Predictive monitoring of traffic rules," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2021, pp. 915–922.
- [37] S. Maierhofer, P. Moosbrugger, and M. Althoff, "Formalization of intersection traffic rules in temporal logic," in *2022 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2022, pp. 1135–1144.
- [38] X. Qin, Y. Xian, A. Zutshi, C. Fan, and J. V. Deshmukh, "Statistical verification of cyber-physical systems using surrogate models and conformal inference," in *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPs)*, 2022, pp. 116–126.
- [39] H. L. Younes, "Probabilistic verification for "black-box" systems," in *International Conference on Computer Aided Verification*. Springer, 2005, pp. 253–265.
- [40] M. Arief, Z. Huang, G. K. S. Kumar, Y. Bai, S. He, W. Ding, H. Lam, and D. Zhao, "Deep probabilistic accelerated evaluation: A robust certifiable rare-event simulation methodology for black-box safety-critical systems," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2021, pp. 595–603.
- [41] C. Harper, G. Chance, A. Ghobrial, S. Alam, T. Pipe, and K. Eder, "Safety validation of autonomous vehicles using assertion-based oracles," *arXiv preprint arXiv:2111.04611*, 2021.