Efficient Black-box Assessment of Autonomous Vehicle Safety

Justin Norden* Trustworthy AI justin@trustworthy.ai Matthew O'Kelly* Trustworthy AI matt@trustworthy.ai Aman Sinha* Trustworthy AI aman@trustworthy.ai

Abstract

While autonomous vehicle (AV) technology has shown substantial progress, we still lack tools for rigorous and scalable testing. Real-world testing, the *de-facto* evaluation method, is dangerous to the public. Moreover, due to the rare nature of failures, billions of miles of driving are needed to statistically validate performance claims. Thus, the industry has largely turned to simulation to evaluate AV systems. However, having a simulation stack alone is not a solution. A simulation testing framework needs to prioritize which scenarios to run, learn how the chosen scenarios provide coverage of failure modes, and rank failure scenarios in order of importance. We implement a simulation testing framework that evaluates an entire modern AV system as a black box. This framework estimates the probability of accidents under a base distribution governing standard traffic behavior. In order to accelerate rare-event probability evaluation, we efficiently learn to identify and rank failure scenarios via adaptive importance-sampling methods. Using this framework, we conduct the first independent evaluation of a full-stack commercial AV system, Comma AI's OPENPILOT.

1 Introduction

The risks of autonomous vehicle (AV) technologies remain largely unknown. A handful of fatal accidents involving AVs highlight the importance of testing whether the system, as a whole, can safely interact with humans. Unfortunately, the most straightforward way to test AVs—in a real-world environment—requires prohibitive amounts of time due to the rare nature of serious accidents. Specifically, a recent study demonstrates that AVs need to drive "hundreds of millions of miles and, under some scenarios, hundreds of billions of miles to create enough data to clearly demonstrate their safety" [Kalra and Paddock, 2016]. To reduce the time, cost, and dangers of testing in the real world, the field has largely turned to simulation-based testing methods.

Most popular strategies for simulation-based testing are built upon decades of techniques developed for validating traditional software systems. Nevertheless, the complexities of modern AV pipelines push the boundaries of these approaches and draw out significant limitations (see Section 1.1 for specific shortcomings of each technique). We broadly illustrate these limitations by outlining desiderata for an AV-testing framework to enable continuous integration and improvement. The term *policy* is used in the rest of this paper to describe an AV's behavior, which is defined as a mapping from observations to actions. In order to tractably measure and improve performance of an AV policy, a testing framework should have the following features:

• **Safety**: Real-world public road testing puts lives in danger. Testing must begin in a controlled setting (*e.g.* simulation or closed courses).

Machine Learning for Autonomous Driving Workshop at the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

^{*}Equal contribution

• **Efficiency**: Industrial AV systems have regular updates. Safety testing must efficiently find failures and compare to previous performance at the pace of development (*e.g.* daily).

• **Coverage**: Understanding coverage of test scenarios is essential to evaluating the true performance of an AV policy. Finding a one-off adversarial situation amounts to finding a minimum in the landscape of AV performance. In contrast, coverage means finding all high-likelihood failure scenarios, or the area of the performance landscape which is below an acceptable level. This latter quantity is required for accurate estimates of risk.

• **Black-box interaction**: Testing a whole system as a black box maintains its integrity and confidentiality. Integrity ensures we capture as accurately as possible the true risk of an AV policy. Combining tests of individual components does not yield this result, and inserting a scaffolding for testing interfaces between components can fundamentally change the nature of failures. Confidentiality of internal system behavior ensures there is no leakage of proprietary knowledge about the AV policy's design; this is especially important to establish cooperation of AV manufacturers with regulators or third-party testers.

• Adaptivity: Since AV policies have regular updates, testing on a static set of previously-defined scenarios or failure modes results in overfitting and false confidence in safety. Adaptive testing automatically modifies test scenarios as AV policies (and their failure modes) change.

• **Unbiasedness**: Unbiasedness implies that estimates of risk are not systemically warped by priors, previous performance, or artifacts of the testing environment that deviate from true operating conditions.

• **Prioritized results**: For actionable insights, safety testing must rank and prioritize failures. In particular, the relative importance of modes within the long tail of failures should be quantified.

Now we outline our approach to satisfying these requirements. We utilize a probabilistic paradigm which we describe as a *risk-based framework* [O'Kelly et al., 2018]—where the goal is to evaluate the *probability of an accident* under a base distribution representing standard traffic behavior. By assigning learned probabilities to environmental states and agent behaviors, the risk-based framework considers performance of the AV policy under a data-driven model of the world. Accompanying this framework, we develop an efficient method to adaptively find failures, return unbiased estimates of risk, and rank failures according to the learned probabilistic model. We implement this framework in a black-box simulation environment that requires minimal interaction with the policy under test.

With these three components—the risk-based framework, an efficient failure-search algorithm, and a black-box simulation environment—we perform the first independent evaluation of a commercially-available AV policy, Comma AI's OPENPILOT. OPENPILOT is an open-source highway driver-assistance system that includes industry-standard components such as deep-learning based perception and model-predictive control. It has driven more than 10 million miles on public highways [Comma AI, 2019a]. By evaluating our approach on this policy, we demonstrate the viability of our framework as a tool for rigorous risk quantification and scalable continuous integration (CI) of modern AV software.

In Section 2, we formally define our risk-based framework. Since we aim to test policies with reasonable performance, our search for failures is a *rare-event simulation* problem [Asmussen and Glynn, 2007]; we illustrate an adaptive nonparametric importance-sampling approach to search for rare events efficiently and in an unbiased manner. In Section 3, we describe a scalable, distributed simulation system. Notably, OPENPILOT's codebase operates on a bespoke distributed architecture in Python and C++, uses hardware acceleration, and has never been ported to a simulation environment. We describe how this integration was performed allowing deterministic, synchronized execution of the otherwise real-time, asynchronous software. Finally, in Section 4 we perform a case study analyzing failures of OPENPILOT over highway scenarios. Our approach estimates OPENPILOT's failure rate of 1 in 1250 miles by using two orders of magnitude fewer simulations than standard methods.

1.1 Related work

A simplistic approach to using simulation for AV policy testing is through a static set of predetermined challenges (*e.g.* Schram et al. [2013]). However, the repeated use of a small number (relative to the variations present in reality) of test-cases often leads to overfitting; passing these tests gives no guarantees of generalization to real-world performance. "Fuzzing" [Hutchison et al., 2018, Dreossi et al., 2019] attempts to improve generalization performance by adding slight variations to predetermined scenarios. Fuzzing considers coverage of code execution paths as opposed to realism or likelihood of failure modes [Koopman and DeVale, 1999, Wagner and Koopman, 2015, Drozd and Wagner, 2018]. As such, this approach is best applied when the space of allowable perturbations is bounded and tractable to cover exhaustively, such as for a subcomponent of a system. Fuzzing to test AV performance as a whole system is challenging for two reasons. First, exhaustively covering a high-dimensional parameter space requires assumptions that can bias results (see *e.g.* Wagner and Koopman [2015]). Second, the space of allowable perturbations for the whole system is usually unbounded or only loosely bounded by physical/biological constraints; these spaces are best described by a probabilistic distribution that weights the importance of perturbations by their likelihoods. For example, the space of *all physically attainable* human-driving behaviors contains many failure scenarios that are extremely improbable in the real world. Regardless of the code executions they spawn, these improbable failures are much less relevant to understanding safety than higher-likelihood failures. Finally, it is not obvious how to translate performance on a fuzzed testing set to a statistical claim on generalization performance.

Formal verification of an AV policy's "correctness" [Kwiatkowska et al., 2011, Althoff and Dolan, 2014, Seshia et al., 2015, O'Kelly et al., 2016] is a potential candidate to reduce the intractability of empirical validation. However, formal verification is exceedingly difficult because driving policies are subject to crashes caused by other stochastic agents, and thus policies can never be completely safe under all scenarios [Shalev-Shwartz et al., 2017]. Furthermore, verification of systems with deep neural networks is an active research question (see *e.g.* Arora et al. [2014], Cohen et al. [2016], Bartlett et al. [2017], and Sinha et al. [2017]) and the problem is NP-hard for most modern deep-learning architectures with non-smooth activations (*e.g.* ReLUs) [Katz et al., 2017, Tjeng and Tedrake, 2017]. Sometimes, verification of individual subcomponents is tractable, but there is no guarantee that this translates to overall safety when all subsystems operate together. One avenue to reduce intractability is to systematically rule out scenarios where the AV is not at fault [Shalev-Shwartz et al., 2017]. Unfortunately, this is a task subject to logical inconsistency, combinatorial growth in specification complexity, and subjective assignment of fault [O'Kelly et al., 2018].

Falsification [Tuncali et al., 2016]—finding *any* failure scenario in the parameter space—is an obvious relaxation of formal verification. Similar to fuzzing, this approach does not account for failure likelihoods. As such, falsification algorithms often return collections of exotic adversarial scenarios, which do not yield actionable insights into prioritizing system updates. Probabilistic variants of falsification [Lee et al., 2018, Koren et al., 2018] attempt to find the highest-likelihood failures, but they require a white-box Markov decision process (MDP) model of the system; defining this model can be challenging for complex systems, and its behavior may deviate from that of the real system.

2 Proposed approach

Below, we formally define our risk-based framework. This framework prioritizes failures based on likelihood and requires minimal, black-box access to the AV policy under test. Additionally, we illustrate the computational challenges of finding failures and then describe our approach which is efficient, adaptive, and unbiased. Finally, we outline various approaches to building a generative model over behaviors for environmental agents that interact with the AV policy under test.

2.1 Risk-based framework

We let P_0 denote the base distribution that models standard traffic behavior and $X \sim P_0$ be a realization of the simulation (*e.g.* weather conditions and driving policies of other agents). For a continuous objective function $f : \mathcal{X} \to \mathbb{R}$ that measures "safety"—so that low values of f(x) correspond to dangerous scenarios—our goal is to evaluate the probability of a dangerous event

$$p_{\gamma} := \mathbb{P}_0(f(X) < \gamma), \tag{1}$$

for some threshold γ . Examples of this objective function f include simple performance measures such as the (negative) maximum magnitude acceleration or jerk during a simulation. More sophisticated measures include the minimum distance to any other object during a simulation or the minimum time-to-collision (TTC) during a simulation. The objective can even target specific aspects of system performance. For example, to stress-test a perception module, the objective can include the (negative) maximum norm of off-diagonal elements of the perception system's confusion matrix. Importantly, this approach is agnostic to the complexity of the AV policy under test, as we view it as a black box: given simulation parameters X, the black box evaluates f(X) by running a simulation. In particular, analysis of systems with deep-learning based perception systems is straightforward in this paradigm, whereas such systems render formal verification methods intractable. Moreover, this framework requires the output of only a single floating point number f(X) from a simulation. This means that neither the simulation nor ego-policy codebases needs to be modified for testing apart from changes already present to run the policy in a simulated environment rather than on a real vehicle. Furthermore, potentially sensitive information about the precise meaning of X within the simulation environment and the AV's decision-making behavior remains confidential from the testing algorithm. One limitation of this approach is the inherent bias introduced by the single objective f. To alleviate this shortcoming, one can also consider extensions that evaluate multiple risk metrics or a stochastic framework where f itself is sampled from a probabilistic distribution. The latter is particularly useful for nondeterministic AV policies. Of course, these extensions also leak more information about the ego policy.

Inefficiency of naive Monte Carlo in rare-event simulation For AV policies which approach or exceed human-level performance, adverse events are rare and the probability p_{γ} is close to 0. Thus, estimating p_{γ} is a rare-event simulation problem (see Asmussen and Glynn [2007] for an overview of this topic). For rare probabilities p_{γ} , the naive Monte Carlo method is prohibitively inefficient. Specifically, for a sample $X_i \stackrel{\text{i.i.d.}}{\sim} P_0$, naive Monte Carlo uses $\hat{p}_{N,\gamma} := \frac{1}{N} \sum_{i=1}^N \mathbf{1} \{f(X_i) < \gamma\}$ as an estimate of p_{γ} . When p_{γ} is small, relative accuracy is the appropriate measure of performance. From the central limit theorem, relative accuracy follows

$$\left|\frac{\widehat{p}_{N,\gamma}}{p_{\gamma}} - 1\right| \stackrel{\text{dist}}{=} \sqrt{\frac{1 - p_{\gamma}}{N p_{\gamma}}} \left|Z\right| + o(1/\sqrt{N}) \text{ for } Z \sim \mathsf{N}(0,1).$$

In order to achieve ϵ -relative accuracy, we need $N \gtrsim \frac{1-p_{\gamma}}{p_{\gamma}\epsilon^2}$ simulations. For light-tailed f(X), then $p_{\gamma} \propto \exp(-|\gamma|)$ so that the required sample size for naive Monte Carlo grows exponentially in $|\gamma|$. In the next section, we use adaptive sampling techniques that sample unsafe events more frequently to make the evaluation of p_{γ} tractable.

2.2 Nonparametric importance-sampling approach

To address the shortcomings of naive Monte Carlo in estimating rare event probabilities p_{γ} , we use a multilevel splitting method [Kahn and Harris, 1951, Botev and Kroese, 2008, Guyader et al., 2011, Webb et al., 2018] combined with adaptive importance sampling (AIS) [Asmussen and Glynn, 2007]. Multilevel splitting decomposes the rare-event probability p_{γ} into conditional probabilities with interim threshold levels $\infty =: L_0 > L_1 \dots > L_K := \gamma$,

$$\mathbb{P}_{0}(f(X) < \gamma) = \prod_{k=1}^{K} \mathbb{P}(f(X) < L_{k} | f(X) < L_{k-1}).$$
(2)

This decomposition introduces a product of non-rare probabilities that are easy to estimate individually. Markov chain Monte Carlo (MCMC) is used to estimate each term; obtaining an accurate estimate is tractable as long as consecutive levels are close and the conditional probability is therefore large. Intuitively, the splitting method iteratively steers samples X_i to the rare set $\{X|f(X) < \gamma\}$ through a series of supersets $\{X|f(X) < L_k\}$ with decreasing objective values $L_k > L_{k+1} \ge \gamma$.

Since it is a priori unclear how to choose the levels $\infty =: L_0 > L_1 \dots > L_K := \gamma$, adaptive variants of multilevel splitting choose these levels online. An elementary instantiation of adaptive multilevel splitting (AMS) chooses each intermediate level L_k , $k \leq K - 1$ by fixing $\delta \in (0, 1)$, the fraction of samples that is discarded at each intermediate step (Algorithm 1 in Appendix A). The discarding fraction δ trades off two dueling objectives; for small values, each term in the product (2) is large and hence easy to estimate by MCMC; for large values, the number of total iterations K until convergence is reduced and more samples (δN) at each iteration can be simulated in parallel. See Appendix A for discussion on the benefits of AMS compared to parametric AIS.

The nonparametric nature of AMS allows efficient discovery of disparate failure modes. However, AMS does not generate an efficient sampler over these modes for further investigation. Instead, we learn a generative model of failures from the empirical distribution of failure modes discovered by

AMS. Specifically, we use normalizing flows as our class of generative models.² Normalizing flows transform a base distribution (usually a standard normal distribution) through a series of deterministic, invertible functions. Using the standard change-of-basis transform, the density $\rho(g(x))$ for some invertible map $g : \mathcal{X} \to \mathcal{Y}$ is given by:

$$y = g(x) \implies \rho(y) = \rho(g^{-1}(y)) \cdot |\det J(g^{-1}(y))|,$$

where $J(\cdot)$ denotes the Jacobian. Composing a sequence of such transforms $g_1(x), g_2(g_1(x)) \dots$ gives expressive transformations of the base density. Rezende and Mohamed [2015] describe a modern form of the approach which ensures the Jacobian is triangular, rendering the determinant computation efficient. Since each transform is parameterized by trainable weights, one can learn a distribution with this architecture of composed transformations by maximizing the log-probability of the observed data in the transformed distribution. Further enhancements to this approach can be found in Kingma et al. [2016] and Papamakarios et al. [2017].

2.3 Data-driven generative modeling of agent behaviors for P₀

To implement our risk-based framework, we first learn a base distribution P_0 of nominal traffic behavior. Although this is a challenging open problem, there are a variety of approaches in the literature to do so via imitation learning [Russell, 1998, Ross and Bagnell, 2010, Ross et al., 2011, Ho and Ermon, 2016] on data of real human drivers. Behavior cloning [Bansal et al., 2018] simply treats imitation learning as a standard supervised learning problem, whereas reinforcement-learning techniques such as generative adversarial imitation learning [Ho and Ermon, 2016] improve generalization performance [Ross and Bagnell, 2010, Ross et al., 2011]. Recent advances in imitationlearning techniques specifically for driving vehicles can be found in Kuefler and Kochenderfer [2018], Codevilla et al. [2018], Rhinehart et al. [2018], Codevilla et al. [2019], and Rhinehart et al. [2019]. In this paper, we use an agent modeling framework previously validated on highway scenarios in the NGSim dataset [US Department of Transportation – FHWA, 2008] by Kuefler et al. [2017] and O'Kelly et al. [2018].

3 Systems

In this section we describe the integration of a widely-deployed autonomous vehicle system, OPEN-PILOT [Comma AI, 2019b], and a popular simulation environment, CARLA [Dosovitskiy et al., 2017], with the theoretical framework of the previous section.³

Creating a test-harness for OPENPILOT OPENPILOT is a SAE Level 2 [SAE, 2014] driverassistance system designed to be integrated with Comma AI's proprietary camera hardware and a customized software platform based on the Android operating system. At the time of writing, OPEN-PILOT has been installed on more than 4500 vehicles and driven more than 10 million miles [Comma AI, 2019a]. A lack of publicly-discussed incidents suggests that it is a capable system, but no rigorous statistics have been released regarding its safety or performance.

Unlike end-to-end systems, OPENPILOT contains a relatively standard set of modules for perception, planning, and control. Figure 1 gives a high-level overview of the four primary components of OPENPILOT: (1) State information is collected from the target vehicle's steering, throttle, and brake modules as well as processed-radar outputs. (2) The EON dash camera contains a camera, processor, and an interface to the vehicle. (3) The collected sensor observations, including images from the camera, are sent to the perception system where camera calibration is performed and a model of the driving environment is created. (4) Finally, the driving environment data is fed to a model-predictive control system which outputs a trajectory to be tracked by a low-level PID controller.

Despite public availability of the source code, the majority of the system has not previously been compiled for x86 hardware, even, to the best of our knowledge, by Comma AI. We develop TEST-PILOT, a minimally invasive fork of OPENPILOT which enables deterministic execution of all components on x86 hardware with support for simulator-time synchronization and optional GPU acceleration within a distributed, containerized environment (see Appendix B for further details).

²As with any importance sampling technique, estimates of risk from our sampler are unbiased.

³The authors of this paper are not affiliated with Comma AI.



Figure 1: Components of OPENPILOT.

There are several differences between the experiments conducted using TESTPILOT below and what a user might experience in an OPENPILOT-equipped vehicle. First, OPENPILOT contains a drivermonitoring system that disengages when a user diverts attention from the road. We do not model these disengagement events in our experiments. Second, most OPENPILOT-equipped vehicles utilize a radar system to aid vision-based vehicle detection. At the time of writing, open-source vehicle simulators such as CARLA do not contain a radar sensor model. Thus, it is possible that a radar-equipped OPENPILOT vehicle may be able to avoid some of the crashes reported in Section 4.

4 Case study

We consider a highway scenario consisting of six agents, five of which are part of the environment. All vehicles are constrained to start within a set of possible initial configurations consisting of pose and velocity, and each vehicle has a goal of reaching the end of an approximately 1 kilometer stretch of road within CARLA's built-in "Town04" environment [Dosovitskiy et al., 2017]. We also modify continuous weather parameters as well as the behaviors of the environment agents. See Appendix C for a full description of the distribution P_0 . We cap simulations at 10 seconds in length (and they may be shorter if there is a collision). Throughout these experiments, we employ minimum timeto-collision (TTC) over a simulation (see Sontges et al. [2018] and Appendix D for a more detailed discussion) as our objective function f(X). First, we compare the efficiency of our approach with naive Monte Carlo in generating statistically-confident estimates of rare-event probabilities. We then characterize the failure scenarios and their implications for OPENPILOT's policy.

4.1 Comparison with naive Monte Carlo

We begin by comparing the output of AMS (without the flow-based importance sampler) to naive Monte Carlo. Namely, we run AMS with a given threshold TTC value γ and allow naive Monte Carlo to use same number of resulting samples (*e.g.* evaluations of the simulator). Figure 2 compares performance of the two algorithms. In Figure 2(a), we show confidence intervals for the estimated probabilities along with a ground-truth value estimated using naive Monte Carlo with roughly 160,000 samples. In Figure 2(b), we plot the ratio of the variance of the two estimators. For events with probability lower than 10^{-3} ($\gamma \leq 0.8$ seconds), AMS outperforms naive Monte Carlo, and the variance of the failure probability is reduced by up to $10 \times$ (when $\gamma \leq 0.1$ seconds). On the other hand, naive Monte Carlo is more efficient than AMS at estimating non-rare events ($\gamma > 0.8$ seconds). Interestingly, for this scenario and metric f, there is a steep exponential drop in likelihood for events near $\gamma \approx 1$ second, which happens to correspond with the regime below which our adaptive methods outperform naive Monte Carlo.

Figure 3 shows results from the flow-based importance sampler built upon the output of AMS. Importantly, to compare the performance of this sampler properly with that of naive Monte Carlo, we allow naive Monte Carlo to have a number of samples equal to the amount sampled by the flow *in addition* to the amount required to run AMS. In Figure 3(a), we illustrate the fact that the flow-based importance sampler places high likelihood in failure regions, sampling three orders of magnitude more failures than naive Monte Carlo for very dangerous events ($\gamma \leq 0.1$ seconds). In Figure 3(b) we see that the importance-sampling estimate has variance two orders of magnitude smaller than that of naive Monte Carlo for events with $\gamma \leq 1$ second. In other words, naive Monte Carlo requires two orders of magnitude more simulations than our approach to generate estimates of failure probabilities with the same level of statistical confidence.

4.2 Analysis of failure probability and failure modes

Failure probability We choose $\gamma = 0.1$ seconds as our threshold for failure, so that we are interested in the probability of events with TTC < 0.1 seconds. This threshold encompasses both



Figure 2. (a) 95% confidence intervals for \hat{p}_{γ} for 7 runs of AMS with 450 initial samples and naive Monte Carlo with the same number of simulation evaluations. 'True' corresponds to naive Monte Carlo with 160,000 samples. (b) Ratio of the variance for naive Monte Carlo and AMS. As expected, AMS boosts in performance increase with rarity (γ smaller).



Figure 3. Ratios for event frequency and variance of \hat{p}_{γ} for flow-based importance sampler built upon AMS output vs. naive Monte Carlo sampling. The naive sampling budget is the number of samples drawn by the flow-based importance sampler *plus* the total number of samples required by AMS.

situations where there is no actual crash but the event is still extremely dangerous as well as situations that result in actual impact. As illustrated in Figure 2(a), the failure rate is roughly 1 in 10^4 simulations, which, given the average speeds of the vehicles (roughly 20 meters per second) and the simulation length (10 seconds), corresponds to 1 in 1250 miles.

The only reference values with which to compare this estimated failure rate are the reported disengagement rates by AV manufacturers, but we do so with strong caveats. Disengagements are defined as situations where a human safety driver takes control of a vehicle that was operating autonomously. Importantly, there is no accepted standard for when a disengagement should occur nor do AV manufacturers publish their internal definitions. In 2018, Comma AI claimed a rate of 1 disengagement per hour of driving [Bigelow, 2018]. Other manufacturers reported a slightly more transparent metric: Waymo claimed a disengagement rate of 1 per 11000 miles, Cruise claimed 1 per 5200 miles, Zoox claimed 1 per 1900 miles, and Nuro claimed 1 per 1000 miles [Crowe, 2018].

Failure modes Since we have access to the OPENPILOT codebase, we can investigate the failure modes. An elementary linear dimensionality reduction of the failure conditions X into 2 dimensions (via PCA) yields 4 visible clusters, as depicted in Figure 4(a). Whereas the first PCA mode largely characterizes changes in weather conditions (*e.g.* sunlight in clusters 1 and 3 vs. rain in clusters 2 and 4), the second PCA mode characterizes changes in relative velocities between vehicles (low relative velocity in clusters 1 and 2 vs. high relative velocity in clusters 3 and 4). The predominant culprit of failures is the vision system; it has high uncertainty about vehicles ahead of it as well as lane boundaries when there is direct sunlight on the camera or heavy rain (Figures 4(b)-(c) and 4(d)-(e)). The crashes are either glancing collisions when relative velocities are low or high-impulse collisions when relative velocities are low or high-impulse collisions of these failures for potential improvements in OPENPILOT are clear: (a) the vision system can be improved to better handle these adverse weather conditions, and (b) the feedback loop between the vision and the control stacks can be made more conservative in the presence of high uncertainty.



(a) Failure clusters (b) Sun - Dashcam (c) Sun - Overhead (d) Rain - Dashcam (e) Rain - Overhead

Figure 4. Analysis of failure modes. (a) Graphical depiction of failure modes in 2 dimensions via PCA. (b)-(c) Views from a crash simulation where direct sunlight caused errors in the deep-learningbased vision system, and OPENPILOT failed to see the blue lead car and rear-ended it. (d)-(e) Views from a crash simulation where rain caused errors in the vision system and OPENPILOT misjudged lane boundaries. A playlist of failure mode examples can be viewed at https://www.youtube.com/ playlist?list=PL4lTzImUFh3xhx7kvjrIh6EZudg2n5R3C.

5 Conclusion

Our risk-based framework, combined with efficient adaptive search in black-box simulation, is an approach to efficiently quantify risk and achieve scalable CI for AVs. Real-world testing is still the (slow, expensive, and dangerous) gold standard for asserting a given level of safety for an AV policy, and it is best applied to confirm rather than replace evidence from simulations. Verified subsystems also provide evidence that components meet their design specifications. However, verification is limited by computational intractability, the quality of the specification, and access to white-box models. Rule-based systems for evaluating failures do not yield tangible or actionable safety results, whereas risk-based models for evaluating AV policies are both more objective and actionable. Additionally, while obvious adversarial "corner cases" are sometimes evocative, the failures that need prioritization are those in the seemingly benign long-tail of less-exotic but more-frequent failure situations. Automating the discovery of these less-exotic yet more-frequent failures is paramount to understanding, deploying, and improving AV systems beyond human-level performance.

The risk-based framework is most useful when the base distribution P_0 is accurate, and otherwise it can suffer from systemic biases. Namely, the unbiasedness of the search algorithm is conditioned on the search space—no search algorithm can inherently correct for a poor simulation environment that has an unrealistic rendering engine or unrealistic models of human behavior. Nevertheless, although an estimate of p_{γ} suffers from bias when P_0 is misspecified, our adaptive sampling techniques still efficiently identify dangerous scenarios which are independent of potentially subjective assignments of blame. Principled techniques for building and validating the model of the environment P_0 represent open research questions.

Importantly, the risk-based framework can easily be combined with existing approaches to testing in simulation. The family of nonparametric importance-sampling algorithms outlined in Section 2.2 is a drag-and-drop replacement for existing methods that perform grid search, fuzzing, or naive Monte Carlo search over a parameter space. Furthermore, metrics that assign blame, such as those outlined by Shalev-Shwartz et al. [2017], can be incorporated into the objective function f and used to guide the search for failures. Finally, even though it is designed to work with the whole blackbox system, our risk-based framework can be applied to search for failures in subsystems (*e.g.* the post-perception planning module) as well.

Overall, our framework offers significant speedups over real-world testing, allows efficient, automated, and unbiased analysis of AV behavior, and ultimately provides a powerful tool to make safety analysis a tractable component in AV design. We believe that rigorous safety evaluation of AVs necessitates adaptive benchmarks that maintain system integrity and prioritize system failures. The framework presented in this paper accomplishes these goals by utilizing adaptive importancesampling methods which require only black-box access to the policy and testing environment. As demonstrated on Comma AI's OPENPILOT, our approach provides a principled, statisticallygrounded method to continuously improve AV software. More broadly, we believe this methodology can enable the feedback loop between AV manufacturers, regulators, and insurers that is required to make AVs a reality.

References

- M. Althoff and J. Dolan. Online verification of automated road vehicles using reachability analysis. *Robotics, IEEE Transactions on*, 30(4):903–918, Aug 2014. ISSN 1552-3098. doi: 10.1109/TRO. 2014.2312453.
- S. Arora, A. Bhaskara, R. Ge, and T. Ma. Provable bounds for learning some deep representations. In *International Conference on Machine Learning*, pages 584–592., 2014.
- S. Asmussen and P. W. Glynn. Stochastic Simulation: Algorithms and Analysis. Springer, 2007.
- M. Bansal, A. Krizhevsky, and A. Ogale. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.
- P. L. Bartlett, D. J. Foster, and M. J. Telgarsky. Spectrally-normalized margin bounds for neural networks. In Advances in Neural Information Processing Systems, pages 6241–6250., 2017.
- P. Bigelow. After a pause, comma.ai delivers a driver-assist system that rivals super cruise and autopilot, July 2018. URL https://www.caranddriver.com/news/a22480267/after-apause-commaai-delivers-a-driver-assist-system-that-rivals-super-cruiseand-autopilot/.
- Z. I. Botev and D. P. Kroese. An efficient algorithm for rare-event probability estimation, combinatorial optimization, and counting. *Methodology and Computing in Applied Probability*, 10(4): 471–505, 2008.
- C.-E. Bréhier, T. Lelièvre, and M. Rousset. Analysis of adaptive multilevel splitting algorithms in an idealized case. *ESAIM: Probability and Statistics*, 19:361–394, 2015.
- F. Codevilla, M. Miiller, A. López, V. Koltun, and A. Dosovitskiy. End-to-end driving via conditional imitation learning. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 1–9. IEEE, 2018.
- F. Codevilla, E. Santana, A. M. López, and A. Gaidon. Exploring the limitations of behavior cloning for autonomous driving. arXiv preprint arXiv:1904.08980, 2019.
- N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on Learning Theory*, pages 698–728., 2016.
- Comma AI, May 2019a. URL https://twitter.com/comma_ai/status/ 1123987351276068864.
- Comma AI. Openpilot. https://github.com/commaai/openpilot, 2019b.
- S. Crowe. Waymo autonomous vehicles leave apple in the dust, February 2018. URL https: //www.therobotreport.com/waymo-autonomous-vehicles-apple/.
- A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. Carla: An open urban driving simulator. arXiv preprint arXiv:1711.03938, 2017.
- T. Dreossi, D. J. Fremont, S. Ghosh, E. Kim, H. Ravanbakhsh, M. Vazquez-Chanlatte, and S. A. Seshia. Verifai: A toolkit for the design and analysis of artificial intelligence-based systems. *arXiv preprint arXiv:1902.04245*, 2019.
- W. Drozd and M. D. Wagner. Fuzzergym: A competitive framework for fuzzing and learning. *arXiv* preprint arXiv:1807.07490, 2018.
- Google. grpc: A high performance, open-source universal rpc framework, 2019. URL https://grpc.io/.
- A. Guyader, N. Hengartner, and E. Matzner-Løber. Simulation and estimation of extreme quantiles and extreme probabilities. *Applied Mathematics & Optimization*, 64(2):171–196, 2011.
- J. Ho and S. Ermon. Generative adversarial imitation learning. In Advances in Neural Information Processing Systems, pages 4565–4573, 2016.

- C. Hutchison, M. Zizyte, P. E. Lanigan, D. Guttendorf, M. Wagner, C. L. Goues, and P. Koopman. Robustness testing of autonomy software. In *Proceedings of the 40th International Conference* on Software Engineering: Software Engineering in Practice, pages 276–285. ACM, 2018.
- A. Ignatov, R. Timofte, W. Chou, K. Wang, M. Wu, T. Hartley, and L. Van Gool. Ai benchmark: Running deep neural networks on android smartphones. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 0–0, 2018.
- H. Kahn and T. Harris. Estimation of particle transmission by random sampling. 1951.
- 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*, 94: 182–193, 2016.
- G. Katz, C. Barrett, D. Dill, K. Julian, and M. Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. *arXiv:1702.01135 [cs.AI]*, 1:1, 2017.
- A. Kendall, M. Grimes, and R. Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE international conference on computer vision*, pages 2938–2946, 2015.
- D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing* systems, pages 4743–4751, 2016.
- P. Koopman and J. DeVale. Comparing the robustness of posix operating systems. In *Digest of Papers. Twenty-Ninth Annual International Symposium on Fault-Tolerant Computing (Cat. No. 99CB36352)*, pages 30–37. IEEE, 1999.
- M. Koren, S. Alsaif, R. Lee, and M. J. Kochenderfer. Adaptive stress testing for autonomous vehicles. In 2018 IEEE Intelligent Vehicles Symposium (IV), pages 1–7. IEEE, 2018.
- A. Kuefler and M. J. Kochenderfer. Burn-in demonstrations for multi-modal imitation learning. In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, pages 1071–1078. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer. Imitating driver behavior with generative adversarial networks. In *Intelligent Vehicles Symposium (IV)*, 2017 IEEE, pages 204–211. IEEE, 2017.
- M. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: Verification of probabilistic real-time systems. In *International conference on computer aided verification*, pages 585–591. Springer, 2011.
- R. Lee, O. J. Mengshoel, A. Saksena, R. Gardner, D. Genin, J. Silbermann, M. Owen, and M. J. Kochenderfer. Adaptive stress testing: Finding failure events with reinforcement learning. *arXiv* preprint arXiv:1811.02188, 2018.
- M. O'Kelly, H. Abbas, S. Gao, S. Shiraishi, S. Kato, and R. Mangharam. Apex: Autonomous vehicle plan verification and execution. volume 1, Apr 2016.
- 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*, pages 9827–9838, 2018.
- G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.
- D. J. Rezende and S. Mohamed. Variational inference with normalizing flows. In Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37, pages 1530–1538. JMLR. org, 2015.
- N. Rhinehart, R. McAllister, and S. Levine. Deep imitative models for flexible inference, planning, and control. *arXiv preprint arXiv:1810.06544*, 2018.

- N. Rhinehart, R. McAllister, K. Kitani, and S. Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. arXiv preprint arXiv:1905.01296, 2019.
- S. Ross and D. Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668, 2010.
- S. Ross, G. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635, 2011.
- R. Y. Rubinstein and D. P. Kroese. The cross-entropy method: A unified approach to Monte Carlo simulation, randomized optimization and machine learning. Information Science & Statistics, Springer Verlag, NY, 2004.
- S. Russell. Learning agents for uncertain environments. In *Proceedings of the eleventh annual conference on Computational learning theory*, pages 101–103. ACM, 1998.
- J. SAE. 3016: 2014 taxonomy and definitions for terms related to on-road motor vehicle automated driving systems. *Society of Automotive Engineers*, 2014.
- R. Schram, A. Williams, and M. van Ratingen. Implementation of autonomous emergency braking (aeb), the next step in euro ncap's safety assessment. *ESV, Seoul*, 2013.
- S. A. Seshia, D. Sadigh, and S. S. Sastry. Formal methods for semi-autonomous driving. In *Proceedings of the 52nd Annual Design Automation Conference*, page 148. ACM, 2015.
- S. Shalev-Shwartz, S. Shammah, and A. Shashua. On a formal model of safe and scalable selfdriving cars. *arXiv preprint arXiv:1708.06374*, 2017.
- A. Sinha, H. Namkoong, and J. Duchi. Certifiable distributional robustness with principled adversarial training. arXiv preprint arXiv:1710.10571, 2017.
- S. Sontges, M. Koschi, and M. Althoff. Worst-case analysis of the time-to-react using reachable sets. In 2018 IEEE Intelligent Vehicles Symposium (IV), pages 1891–1897. IEEE, 2018.
- V. Tjeng and R. Tedrake. Verifying neural networks with mixed integer programming. *arXiv:1711.07356 [cs.LG]*, 2017.
- C. E. Tuncali, T. P. Pavlic, and G. Fainekos. Utilizing s-taliro as an automatic test generation framework for autonomous vehicles. In *Intelligent Transportation Systems (ITSC)*, 2016 IEEE 19th International Conference on, pages 1470–1475. IEEE, 2016.
- US Department of Transportation FHWA. Ngsim next generation simulation, 2008.
- M. Wagner and P. Koopman. A philosophy for developing trust in self-driving cars. In *Road Vehicle Automation 2*, pages 163–171. Springer, 2015.
- S. Webb, T. Rainforth, Y. W. Teh, and M. P. Kumar. A statistical approach to assessing neural network robustness. 2018.

A Adaptive multilevel splitting

Algorithm 1 Basic adaptive multilevel splitting (AMS)

1: Input: Discarding factor δ , # of MCMC steps T, termination threshold γ , N initial samples $X_i \stackrel{\text{i.i.d.}}{\sim} P_0$ 2: Initialize $L_0 \leftarrow \infty$, $\log(\hat{p}_{\gamma}) \leftarrow 0$, $k \leftarrow 0$ 3: while $L_k > \gamma$ do 4: $k \leftarrow k+1$ 5: Evaluate and sort $f(X_i)$ in decreasing order $L_k \leftarrow \max\{\gamma, f(X_{(\delta N)})\}$ 6: $\hat{P}_k \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbf{1} \{ f(X_i) < L_k \}.$ 7: 8: $\log(\hat{p}_{\gamma}) \leftarrow \log(\hat{p}_{\gamma}) + \log(\hat{P}_k)$ Discard $X_{(1)}, \ldots, X_{(\delta N)}$ and resample with replacement from $X_{(\delta N+1)}, \ldots, X_{(N)}$ 9: Apply T MCMC transitions to each of $X_{(1)}, \ldots, X_{(\delta N)}$ 10: 11: end while

The nonparametric AMS approach complements parametric AIS methods such as the cross-entropy method [Rubinstein and Kroese, 2004]. Parametric AIS methods postulate models for the form of the optimal importance sampling distribution $P_0(\cdot|f(\cdot) < \gamma)$, and they require iterative optimization of objectives with likelihood ratios. Both of these issues can be challenging in high dimensions, wherein failure regions are likely discontinuous and likelihood ratios are numerically unstable. AMS does not require optimization over terms with likelihood ratios nor does it need to postulate models for the failure region(s). On the other hand, the "modes" of failure AMS discovers are limited by the number of samples and the mixing properties of the MCMC sampler employed. Furthermore, contrary to parametric AIS methods, AMS has several convergence guarantees including those for bias, variance, and runtime (see Bréhier et al. [2015] for details). Notably, AMS is unbiased and has relative variance which scales as $\log(1/p_{\gamma})$ as opposed to $1/p_{\gamma}$ for naive Monte Carlo (cf. Section 2.1). Intuitively, AMS computes $O(\log(1/p_{\gamma}))$ independent probabilities, each with variance independent of p_{γ} .

B TESTPILOT implementation

A primary challenge associated with testing OPENPILOT in a simulated environment is the tight integration between the neural networks in the perception system and the EON hardware. The networks are distributed in binary form and use the Snapdragon Neural Processing Engine (SNPE) accelerator [Ignatov et al., 2018]. The first network, "posenet," infers the motion of the camera sensor (rather than to localize the vehicle in a map as in Kendall et al. [2015]). The predicted motion allows OPENPILOT to estimate the three-dimensional locations of objects detected by the "driving_model" network. TESTPILOT utilizes an SNPE SDK feature which enables the execution of network binaries using OpenCL. Additionally, we create a "fake" camera driver which inputs simulated, recorded, or real images from any source into an OpenCL memory buffer.

Deploying the planning and control elements together with the perception pipeline in a simulation environment requires further refinements. The perception and planning elements use a real-time clock to stamp vehicle-state messages. However, the simulation clock does not necessarily match wall-clock time. To address this problem, TESTPILOT supplies OPENPILOT with a clock signal from the simulator. Furthermore, OPENPILOT heavily utilizes asynchronous message passing; this induces non-determinism in the execution order of the control stack. In order to improve the repeatability of the experiments in Section 4, TESTPILOT forces sequential execution of the control components. Lastly, OPENPILOT utilizes a PID controller operating at 100Hz to actuate steering and throttle mechanisms and track the desired state (updated by the planning module at 20Hz). In order to avoid significant time costs and execute the simulator at 20Hz, we send the planning module's desired state to the simulator and perform tracking on the simulator side.

Running simulations at scale To support running simulations at scale, we package OPENPILOT modules and CARLA simulation servers as containerized applications within a standard container orchestration scheme on commodity cloud servers. Our adaptive sampling algorithm coordinates jobs via the open-source remote-procedure-call framework gRPC [Google, 2019]. For some context, we simulated roughly 40,000 miles to generate the results in Section 4 comparing naive Monte Carlo

with our proposed approach. As we show below, our approach requires orders of magnitude fewer simulations to estimate risk with a given precision compared to naive Monte Carlo. A sample of naively drawn simulations can be found here: https://youtu.be/6nuiLkecik8.

C Simulation parameters

The 1 kilometer stretch of road that we consider consists of a straight section followed by a rightturning bend and a final straight section. Our search space P_0 consists of weather parameters, initial positions and velocities of the vehicles, and the behaviors of the environmental agents.

4 parameters of the weather are varied: precipitation on the ground (P_g) , the altitude angle of the sun (A), cloudiness (C), and precipitation in the air (P_a) . Precipitation on the ground follows the distribution $P_g \sim 50$ Uniform(0, 1), where the units are proprietary for the CARLA simulator. Sun altitude follows the distribution $A \sim 90$ Uniform(0, 1), whose units are in degrees. Letting $C_b \sim \text{Beta}(2, 2)$ and $C_u \sim \text{Uniform}(0, 1)$, cloudiness has a distribution given by the following mixture:

 $C \sim 30C_b \mathbf{1}\{C_u < 0.5\} + (40C_b + 60)\mathbf{1}\{C_u \ge 0.5\},\$

where $\mathbf{1}\{\cdot\}$ is the indicator function and the units of C are in proprietary units for CARLA. Precipitation in the air is a deterministic function of cloudiness, $P_a = C\mathbf{1}\{C \ge 70\}$, where the units are again proprietary CARLA units.

Initial positions for each vehicle are independently given by the pair (S,T), where $S \sim 500\text{Beta}(2,2) + 200$ denotes the longitudinal position along the stretch of road with respect to the start of the road, and $T \sim 0.5\text{Beta}(2,2) - 0.25$ denotes the lateral distance from the lane center with left being positive. Both have units in meters. Initial velocities are independently given by $V \sim 10\text{Beta}(2,2) + 15$ with units in meters/second.

For the behaviors of the environment vehicles we follow the parametrization of O'Kelly et al. [2018]. Namely, we search over the last layer of a neural network GAIL generator model. We use $\xi \in \mathbb{R}^{404}$ to denote the weights of the last layer of the neural network trained to imitate human-like driving behavior with $\xi \sim \mathcal{N}(\mu_0, \Sigma_0)$, with the mean and covariance matrices learned via empirical bootstrap. The input for this GAIL generator is described in O'Kelly et al. [2018]. The single network defined by ξ drives the behavior for all environment agents in a particular simulation rollout.

Together, the weather, position, velocity, and behavior variables form samples $X \in \mathbb{R}^{426}$.

D TTC calculation

Let $T_i(t)$ be the instantaneous time-to-collision between the ego vehicle and the *i*-th environment vehicle at time step t. For a given simulation rollout X, the TTC metric is given by

$$f(X) := \min_{t} \left(\min_{i} T_i(t) \right)$$

The value $T_i(t)$ can be defined in multiple ways [Sontges et al., 2018]. In this study, we define it as the amount of time that would elapse before the two vehicles' bounding boxes intersect assuming that they travel at constant fixed velocities from the snapshot at time t.