

Model-based Validation as Probabilistic Inference

Harrison Delecki

Anthony Corso

Mykel J. Kochenderfer

*Department of Aeronautics and Astronautics
Stanford University, Stanford, CA 94305, USA*

HDELECKI@STANFORD.EDU

ACORSO@STANFORD.EDU

MYKEL@STANFORD.EDU

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

Estimating the distribution over failures is a key step in validating autonomous systems. Existing approaches focus on finding failures for a small range of initial conditions or make restrictive assumptions about the properties of the system under test. We frame estimating the distribution over failure trajectories for sequential systems as Bayesian inference. Our model-based approach represents the distribution over failure trajectories using rollouts of system dynamics and computes trajectory gradients using automatic differentiation. Our approach is demonstrated in an inverted pendulum control system, an autonomous vehicle driving scenario, and a partially observable lunar lander. Sampling is performed using an off-the-shelf implementation of Hamiltonian Monte Carlo with multiple chains to capture multimodality and gradient smoothing for safe trajectories. In all experiments, we observed improvements in sample efficiency and parameter space coverage compared to black-box baseline approaches. This work is open sourced.¹

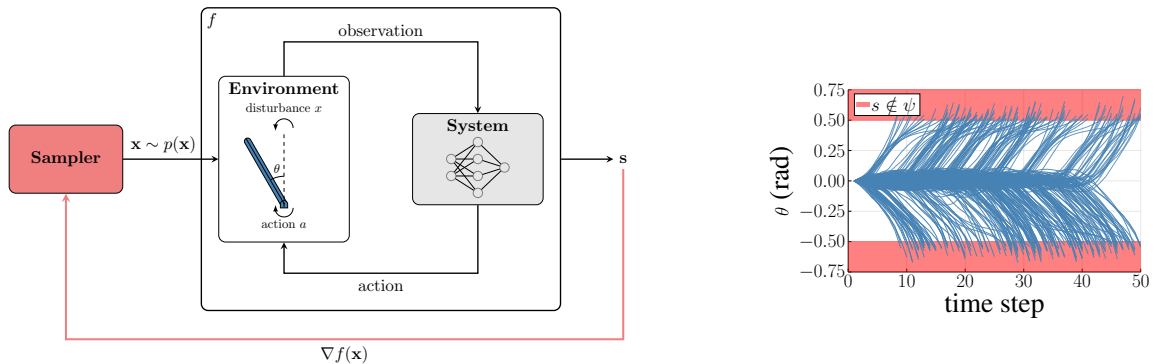
Keywords: safety validation, Bayesian inference

1. Introduction

Greater levels of automation are being considered in safety-critical applications such as self-driving cars (Badue et al., 2021) and air transportation systems (Straubinger et al., 2020). These systems typically operate sequentially by taking actions after observing the environment. Sequential decision-making systems in safety-critical domains require thorough validation for acceptance and safe deployment. Information about potential failure modes is critical for engineers and policy-makers to determine whether certain failure modes require additional engineering attention and to estimate the risk associated with deployment (Corso et al., 2021). Therefore, a key step of safety validation is to determine the distribution over failure trajectories, or search for possible failure modes and determine their associated likelihood.

Searching for failures in safety-critical systems is challenging for several reasons. First, failures tend to be rare. Safety-critical autonomy is designed to be relatively safe from the outset, so naive methods like direct Monte Carlo sampling require an enormous number of samples to discover rare failures. Second, the search space is high dimensional. Autonomous systems typically have large state spaces and operate over long time horizon trajectories. Third, sequential systems can exhibit multimodal failures. Accurately estimating multimodal distributions requires methods that can explore the parameter space without converging to individual modes.

1. <https://github.com/sisl/ModelBasedValidationInference>



(a) Model-based sampling of the distribution over failures.

(b) Sampled failure trajectories of the inverted pendulum.

Figure 1: A diagram of the proposed approach for estimating the distribution over failures. A sampler uses gradients of system dynamics to find disturbances x that lead to a failure, such as toppling an inverted pendulum. The distribution over failures is estimated by drawing many samples that result in failure.

Several previous approaches to safety validation perform falsification (Dreossi et al., 2015; Akazaki et al., 2018; Tuncali and Fainekos, 2019), which aims to find a single failure example or failure mode. The main drawback of these methods is that they tend to converge to a single failure trajectory, such as the most extreme or most likely example (Lee et al., 2020). The mode-seeking behavior of falsification methods is insufficient to capture the distribution over failure trajectories. Importance sampling (Kim and Kochenderfer, 2016; Zhao et al., 2017; O’Kelly et al., 2018) and Markov Chain Monte Carlo (Norden et al., 2019; Sinha et al., 2020) approaches have also been used to sample failure trajectories. These approaches aim to approximate the distribution over failures using samples. However, existing approaches are only applicable in low-dimensional spaces and rely on ad-hoc algorithmic modifications. These limitations make existing approaches less useful for the validation of sequential systems in a general setting.

In this work, we aim to enable efficient approximation of the distribution over failure trajectories in sequential systems. First, we frame the problem as Bayesian inference of the distribution over failures due to disturbances in the environment. We then propose a modeling framework to represent the failure distribution in a probabilistic programming paradigm. This model-based representation enables us to easily use advanced inference algorithms for high-dimensional sampling problems, such as Hamiltonian Monte Carlo (Neal, 2011; Hoffman et al., 2014). The proposed approach is illustrated in Fig. 1. Our contributions are as follows:

- We frame sampling the distribution over failures in sequential systems as Bayesian inference.
- We propose a model-based approach to sample from the distribution over failures.
- We demonstrate the proposed approach on three sample validation problems.

2. Related Work

Relevant related work can be grouped in three categories: white-box validation methods, falsification approaches, and sampling-based approaches.

White-box methods: Traditional methods in system validation often focus on white-box techniques, which assume that a full mathematical model of the system is known (Katoen, 2016; Clarke et al., 2018). These methods typically do not scale to large problems because they consider every possible execution of the system and the environment (Alur, 2015). The goal of this work is similar in that we also use knowledge internal to the system to search for multiple possible failure events.

Falsification: Several recently proposed safety validation methods for sequential systems take a falsification approach. These methods use optimization (Deshmukh et al., 2017), trajectory planning (Dreossi et al., 2015; Tuncali and Fainekos, 2019), or reinforcement learning (Akazaki et al., 2018) to find disturbances in the environment that lead to any failure or the most severe failure. Adaptive Stress Testing is a reinforcement learning method that aims to find the most-likely failures according to a prescribed probability model (Lee et al., 2015; Koren et al., 2018). Falsification methods tend to converge to a single failure mode rather than exploring the distribution over failures.

Sampling-based approaches: Importance sampling has been applied to estimate the distribution over failures for aircraft collision avoidance systems and autonomous vehicles (Kim and Kochenderfer, 2016; O’Kelly et al., 2018). The cross-entropy method (CEM) aims to iteratively learn the optimal importance sampling distribution. CEM relies on the specification of a parameterized family of distributions and may fail when the true underlying distribution is multimodal or high dimensional (Geyer et al., 2019). Markov Chain Monte Carlo (MCMC) has also been applied to estimate the distribution over failures (Botev and Kroese, 2008; Norden et al., 2019). MCMC methods used in previous work explore the distribution using a random walk or hand-crafted proposal distributions (Neal, 2011). While this is effective for lower dimensional problems, this exploration strategy struggles to scale to high dimensional distributions. In this work, we use Hamiltonian Monte Carlo to estimate the distribution over failure trajectories. HMC is an MCMC algorithm that uses gradients rather than a random-walk or a hand-crafted proposal to explore the distribution.

3. Approach

In this section, we present our framework for model-based Bayesian inference of the failure distribution in sequential systems. First, we introduce the necessary notation and definitions. Then, we describe our approach to estimate the distribution over failures through probabilistic inference.

3.1. Formulation

We build upon the notation of Corso et al. (2021) to describe the safety validation problem. Consider a system under test that takes actions in an environment after receiving observations. We define a safety property ψ that we aim to evaluate for the system under test. The property ψ is defined over the state trajectories of the environment, $\mathbf{s} = [s_1, \dots, s_t]$, where s_t is the state of the environment at time t . If the trajectory results in a failure, we say that $\mathbf{s} \notin \psi$.

We perturb the environment with disturbances with the goal of inducing the system under test to violate the safety specification. A disturbance trajectory \mathbf{x} has probability density $p(\mathbf{x})$. We assume that disturbances determine all sources of stochasticity in the environment. For example, sensor noise and stochastic dynamics are potential disturbances that could lead to failure.

We denote the simulation of the system under test with environment disturbances by a dynamics function f . Note that f represents the dynamics of both the system under test and the environment. The resulting state trajectory under disturbances is written as $\mathbf{s} = f(\mathbf{x})$. Since all stochasticity in the simulation is determined by \mathbf{x} , the disturbance model $p(\mathbf{x})$ induces a distribution over state

trajectories, $p(\mathbf{s})$. We seek the distribution over state trajectories that violate the safety property, the conditional distribution $p(\mathbf{s} \mid \mathbf{s} \notin \psi)$.

3.2. Approach Overview

Environment dynamics or characteristics of the system under test often make it impossible to express the failure distribution analytically. However, it is usually easy to sample from these systems by simulating the system under test with disturbances. Bayesian inference methods such as MCMC approximate a target distribution by drawing samples. In this work, we frame the safety validation problem as Bayesian inference of the distribution over failures.

Sampling-based Bayesian inference requires a way to draw samples from a target distribution. We represent the state trajectory distribution as a probabilistic model using probabilistic programming (Gordon et al., 2014). Under this paradigm, we can describe the distribution over state trajectories as a function of random variables that represent the disturbance model. Modern probabilistic programming languages, such as Turing (Ge et al., 2018), let us wrap existing simulations in a probabilistic program with little modification of the underlying simulation.

This probabilistic model represents the distribution over state trajectories induced by the disturbance model. We represent the distribution over failures by conditioning the probabilistic model on the desired outcome of the simulation — that the simulated trajectory is a failure. Any safe trajectory has zero likelihood under the conditional distribution. We represent this condition using the Dirac delta distribution, whose value is zero for all state trajectories except when $\mathbf{s} \notin \psi$. Under this formulation, the log-likelihood of a state trajectory $\ell(\mathbf{s} \mid \mathbf{s} \notin \psi)$ is

$$\ell(\mathbf{s} \mid \mathbf{s} \notin \psi) = \log \delta(\mathbf{s} \notin \psi) + \log p(\mathbf{s}) \quad (1)$$

The main challenges associated with sampling from $p(\mathbf{s} \mid \mathbf{s} \notin \psi)$ are the distribution’s high dimensionality, discontinuity, and multimodality. The following sections show how we address each of these challenges.

3.3. Model-based Approach with Automatic Differentiation

State trajectories in sequential systems of interest tend to be high dimensional. It is well established that black-box, gradient-free sampling from high-dimensional distributions becomes difficult with increasing dimensionality (Neal, 2011). Gradient-based sampling algorithms such as Hamiltonian Monte Carlo (HMC) (Duane et al., 1987; Girolami and Calderhead, 2011) use information about the shape of the underlying distribution to take large steps in parameter space, allowing them to scale to higher dimensions.

To address the challenge of sampling high-dimensional failure trajectories in sequential systems, we take a model-based approach. Gradients of the dynamics function f with respect to disturbances are computed using automatic differentiation. The simulation environment and system under test are modeled in a framework compatible with automatic differentiation (AD). AD enables algorithms to compute exact gradients of the posterior likelihood density with respect to the input disturbances. We use Zygote (Innes et al., 2019), which performs source-to-source AD of Julia code.

In contrast with some previous approaches that assume the system is Markov (Corso et al., 2020), our approach can handle non-Markov systems. We estimate the distribution over failures using rollouts of the system dynamics under disturbances, which fully characterize the system behavior and allows us to validate policies that depend on state estimates rather than true states.

3.4. Smoothing Approximation for Sharp Gradients

Gradients of the posterior density will be undefined in regions where the state trajectory is safe due to the Dirac delta in Eq. (1). Additionally, the boundary between safe and failing state trajectories is discontinuous. As a result, gradient-based sampling algorithms have no information about how to improve a proposal. We relax the discontinuity by approximating the Dirac delta distribution as a one-dimensional Gaussian with a small variance ϵ . We also assume that we can calculate a distance to failure $\Delta(s)$ for any safe trajectory. The likelihood of the distance to failure under the Gaussian approximation of the discontinuity is added to the likelihood of the sampled trajectory. The log-likelihood of a trajectory $\ell(s)$ under this smoothing approximation is

$$\ell(s \mid s \notin \psi) = \log \mathcal{N}(\Delta(s) \mid 0, \epsilon) + \log p(s)$$

The role of AD is to calculate $\nabla_x \ell(s \mid s \notin \psi)$, which is now defined for safe and failure trajectories. The Gaussian smooths the discontinuity between the safe and failure regions of the state trajectory space while changing the density in the failure region by a constant. This smoothing approximation along with AD enables us to use gradient-based sampling methods.

The variance ϵ acts as a hyperparameter, with smaller values better approximating the discontinuity and larger values yielding smoother posterior densities. Consider a validation problem for illustrative purposes only, where the state distribution is the unit Gaussian and failure occurs when $|s| > 5$. Fig. 2 shows the log-likelihood of single-step state trajectories for various values of ϵ . As ϵ decreases towards zero, the resulting distribution is a better approximation of the true failure distribution, shown in red. Larger values of ϵ smooth the gap between the separated failure modes on the far left and right tails of the distribution. Empirically, we find values in the range $\epsilon \in [0.01, 0.1]$ work well for the problems investigated.

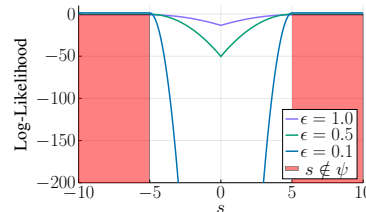


Figure 2: Illustration of smoothing hyperparameter ϵ .

3.5. Sampling for Multimodality

The generality of the approach makes available a wide variety of off-the-shelf Bayesian inference methods. We use HMC with No-U-Turn Sampling (NUTS) (Hoffman et al., 2014), which is a general HMC variant that excels at sampling from complex, high-dimensional posterior distributions.

Multimodality in the distribution over failure trajectories poses a challenge for many MCMC algorithms (Rudoy and Wolfe, 2006). Sampling algorithms may converge to a single mode rather than mixing across the modes in parameter space. In our illustrative example, a single MCMC chain may only sample from the left or right mode of the failure distribution. One way to explore a multimodal space is to run many chains of MCMC from different starting points. In our sampling approach, we use multiple NUTS chains with different starting disturbance trajectories. The number of chains selected depends on the difficulty of the problem, with higher dimension and more modes generally requiring more chains.

4. Experiments

This section presents the experimental evaluation of the proposed approach. First, we discuss our evaluation metrics and baseline sampling approaches. Then, we describe the three simulation environments and policies used in validation experiments.

Metrics and Baselines: We evaluate the proposed approach based on failure rate, log-likelihood of failures discovered, and computation time. The failure rate is the proportion of drawn samples that result in failure. We also use a coverage metric based on the mean dispersion of failure trajectories in disturbance space (Esposito et al., 2004). The mean dispersion $C_{disp} \in [0, 1]$ is computed over a grid in trajectory space as

$$C_{disp} = 1 - \frac{1}{n} \sum_{j=1}^n \frac{\min(d_j, g)}{g}$$

where n is the number of grid points, g is the grid spacing, and d_j is the minimum distance from the grid point to a point in the sampled disturbance trajectories. Larger values of C_{disp} indicate that a sampler has discovered failures in more regions of disturbance trajectory-space.

We compare the proposed gradient-based approach to direct Monte Carlo (MC) sampling and the black-box MCMC Particle Gibbs (PG) (Andrieu et al., 2010) algorithm with 1000 particles.

4.1. Inverted Pendulum

The first system we validate is an inverted pendulum with a neural network controller. While this is not a safety-critical system, it demonstrates the ability of the proposed approach to validate policies with machine learning components. In the inverted pendulum, a policy exerts a control torque to keep the pendulum balanced upright. The state of the system is $s = (\theta, \dot{\theta})$ where θ is the angle of the pendulum from vertical and $\dot{\theta}$ is the angular velocity. The control policy takes continuous-valued actions $a \in [-T_{\max}, T_{\max}]$, where T_{\max} is the maximum control effort of the actuator. The pendulum is underactuated, meaning that the available torque is too weak to result in arbitrary state trajectories. If the pendulum reaches beyond about 0.5 rad from vertical, the available torque will not be able to overcome the torque due to the pendulum’s mass, and the pendulum will fall. The pendulum and failure condition is illustrated in Fig. 3(a).

We optimize a neural network policy using the Proximal Policy Optimization (Schulman et al., 2017) reinforcement learning algorithm, which has been shown to perform well on the inverted pendulum. The policy is represented by a two-layer neural network with 32 neurons per layer and tanh activations between hidden layers. The policy is trained for 10^6 time steps.

We sample failures of control policy by considering external torque disturbances from the environment. Adversarial torque disturbances are assumed to be distributed according to a zero-mean Gaussian with variance σ_n . This experiment uses $\sigma_n = 0.1 \text{ N m}$ and $T_{\max} = 2 \text{ N m}$ so that disturbances are likely to be much smaller than the maximum control torque. A failure occurs when the torque due to gravity overcomes the maximum control torque, which occurs when $|\theta| \geq 0.5 \text{ rad}$. The distance to failure is the minimum distance from the failure threshold. Simulations span a 5 s horizon with a 0.1 s time step. We sample 10,000 samples for each method, with 10 MCMC chains and 1000 samples per chain.

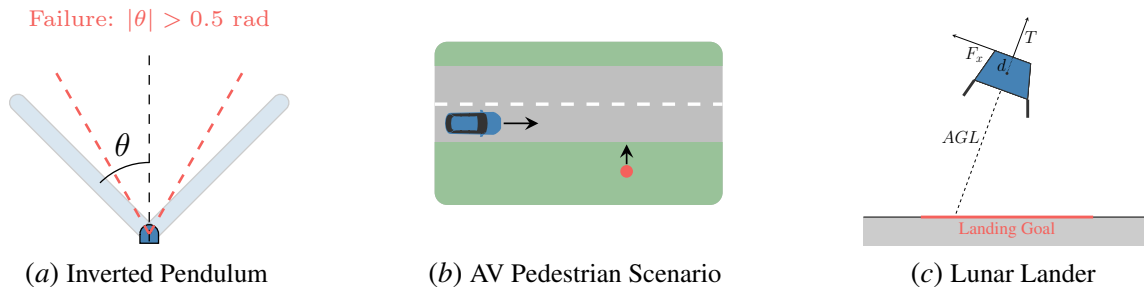


Figure 3: The three environments used in validation experiments.

4.2. Autonomous Vehicle Scenario

The second experiment involves an autonomous vehicle (AV) approaching a single pedestrian at a crosswalk and is based on the work of [Koren et al. \(2018\)](#). This scenario is illustrated in Fig. 3(b). The system under test is the autonomous vehicle’s control policy, which is defined by the Intelligent Driver Model (IDM) ([Treiber et al., 2000](#)). The IDM is a model-based controller that keeps the vehicle in its lane, following the traffic or other obstacles ahead while maintaining a safe distance. The goal of the AV is to identify the pedestrian crossing the road and apply adequate braking to avoid coming too close. Each agent’s state is represented by a 4-tuple $s = (r_x, r_y, v_x, v_y)$ where (r_x, r_y) is the agent’s position and (v_x, v_y) is the agent’s velocity. In the noise-free case, the AV comes to a stop and the pedestrian safely crosses.

The adversary has control over pedestrian acceleration and noise in the AV’s sensors. The disturbances to the system are represented by a 6-tuple $x = (\delta r_x, \delta r_y, \delta v_x, \delta v_y, a_x, a_y)$ where $(\delta r_x, \delta r_y)$ is the noise in the pedestrian’s observed position, $(\delta v_x, \delta v_y)$ is the noise in the pedestrian’s observed velocity, and (a_x, a_y) is the pedestrian’s acceleration. The noise components are processed at each time step and added to the vehicle’s observation of each pedestrian’s position and velocity. The disturbances are assumed to be distributed according to a multivariate Gaussian distribution with zero mean and diagonal covariance matrix. The diagonal components of covariance are $\sigma_r^2 = 0.1$ for position measurement noise, $\sigma_v^2 = 0.1$ for velocity measurement noise, $\sigma_{lat}^2 = 0.01$ over the pedestrian’s lateral acceleration, and $\sigma_{lon}^2 = 0.1$ over the pedestrian’s longitudinal acceleration. These values are chosen to reflect relatively small sensor noise, and that the pedestrian is less likely to accelerate longitudinally while crossing the street.

Failures of this scenario are defined by any collision between the pedestrian and vehicle. The distance to failure is the minimum Euclidean distance between the vehicle and the pedestrian. We sample 20,000 samples for each method, with 20 MCMC chains and 1000 samples per chain.

4.3. Partially Observable Lunar lander

The final validation problem is a simulation of a lunar lander with partially observed state information. The objective of the lunar lander’s policy is to guide the vehicle in a target area with a soft landing. The vehicle state is represented by the 6-tuple $(x, y, \theta, \dot{x}, \dot{y}, \dot{\theta})$, where x is the horizontal position, y is the vertical position, and θ is the vehicle’s orientation. The vehicle’s dynamics are deterministic. The vehicle makes noisy observations of its angular rate ω , horizontal speed v , and above ground level altitude (AGL). AGL is a noisy measurement of the distance from the vehicle’s center of mass to the ground along the vehicle’s longitudinal axis. The continuous actions

are defined by (T, F_x, d) where T is the main thrust acting along the longitudinal axis, F_x is lateral corrective thrust acting at an offset of d from the center of mass. The lander is illustrated in Fig. 3(c).

We optimize a control policy by considering a discrete, fully observable version of the continuous, partially observable problem. We use approximate value iteration to solve for a table-based policy in the discrete problem. Next, we create a differentiable representation of the table-based policy by performing behavior cloning with a neural network approximator. The network is trained to minimize the mean squared error between the continuous output and the table-based action.

We assume that observation noise is identically and independently distributed at each time step according to a zero-mean Gaussian with a diagonal covariance matrix. We use $\sigma_\omega^2 = 0.02$, $\sigma_v^2 = 0.1$, and $\sigma_{AGL}^2 = 1.0$. An extended Kalman filter maintains a multivariate Gaussian belief over the true state in the partially observable problem. The policy uses the mean to calculate an action.

Our approach samples sequences of observation noise that lead to a hard landing, defined as the landing velocity being greater than 6 m/s. Empirically, we discovered the lunar lander contained multiple isolated failure modes. We select initial conditions for the MCMC samplers by performing black-box optimization of the probabilistic model’s posterior density. Specifically, we use particle swarm optimization with five particles and 100 iterations to select initial conditions for each MCMC chain. We sample 30 MCMC chains with 1000 samples each to explore the different failure modes.

5. Results

This section presents experimental results for each validation problem. Probabilistic models of each system were written using Turing and sampled using the HMC-NUTS implementation provided by Ge et al. (2018). Gradients were computed using the Zygote (Innes et al., 2019) AD framework. Experiments were ran in serial on a desktop with 32GB of RAM and an Intel i7-7700K CPU.

Overview: Table 1 shows the failure rate, failure trajectory log-likelihood, sampling time per failure, and mean dispersion for each sampling method and validation setting. The proposed approach consistently outperforms baselines in terms of failure rate. The gradient-based HMC is able to consistently sample from the high-dimensional failure region, while PG and MC samplers struggle. Additionally, the results suggest that HMC consistently samples from the high-likelihood failure region. The computational cost of computing gradients through the state trajectories is outweighed by the increase in sampling performance. HMC samples failures 5–10 times faster in wall clock time compared to baselines. Failures sampled using HMC also cover a significantly greater portion of the disturbance trajectory space across all experiments. Greater disturbance space coverage generally corresponds to a more diverse set of failures being sampled.

Inverted Pendulum: Fig. 4 shows failures sampled using the HMC and baseline methods. The number of failure trajectories for HMC is downsampled for clarity. HMC finds a more diverse set of failure modes compared to PG and MC, with many failures in the positive- and negative- θ direction. The failures found by HMC span a wide range of trajectory lengths and failure modes. HMC uses the gradient of the trajectory with respect to both the pendulum dynamics and the neural network controller, enabling much broader coverage of the failure distribution than black-box baselines.

HMC samples cover the disturbance space well and uncover asymmetric failure modes. The most likely failure trajectories found with each method lead to the pendulum falling in the negative- θ direction. HMC failure samples indicate that the neural network controller exhibits different failure modes in the positive- and negative- θ directions. Positive- θ failures have a lower mean log-likelihood of -28.1 , while negative- θ failures have a mean of -17.4 . The pendulum tends to

Table 1: Results for the inverted pendulum (IP), autonomous vehicle scenario (AV), and partially observable lunar lander (POL). Failure rate, trajectory mean and maximum log-likelihood (LL), sampling time per failure, mean dispersion C_{disp} , and total sampling time are reported for each algorithm. HMC outperforms black-box baselines in terms of failure rate, sampling time per failure, and disturbance-space coverage.

	Method	Failure Rate	Mean LL	Max. LL	Failures/s	C_{disp}	Time (min)
IP	HMC	9.8×10^{-1}	-24.66 ± 9.59	-8.18	8.0	7.5×10^{-1}	22
	PG	2.0×10^{-4}	-21.42 ± 0.15	-19.89	1.1×10^{-3}	2.1×10^{-4}	27
	MC	3.0×10^{-3}	-18.44 ± 2.97	-13.81	1.2	2.4×10^{-2}	2
AV	HMC	7.8×10^{-1}	-2.01 ± 20.00	0.53	7.8	2.0×10^{-2}	40
	PG	1.7×10^{-1}	-0.71 ± 0.48	0.11	6.5×10^{-3}	1.0×10^{-5}	40
	MC	5.0×10^{-4}	-12.05 ± 0.06	-12.35	8.5×10^{-1}	1.0×10^{-9}	3
POL	HMC	1.6×10^{-1}	1.68 ± 1.69	3.19	4.1×10^{-1}	2.1×10^{-1}	375
	PG	2.4×10^{-4}	1.80 ± 0.21	2.12	3.3×10^{-5}	3.2×10^{-3}	810
	MC	3.0×10^{-5}	2.34 ± 1.09	2.61	1.7×10^{-1}	1.5×10^{-5}	13

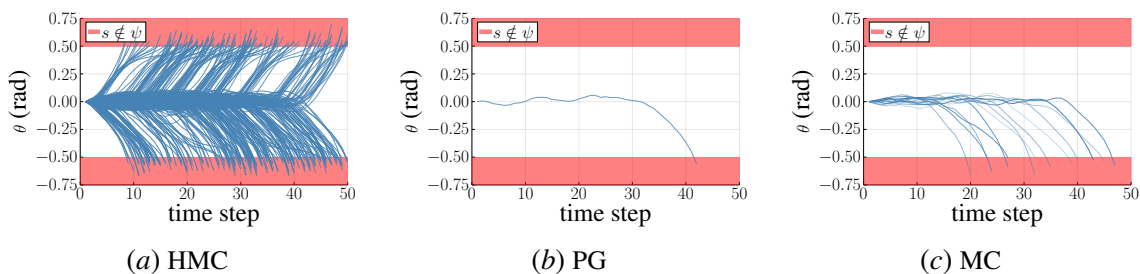


Figure 4: Failure trajectories for the inverted pendulum with HMC and baseline methods. HMC discovers many likely failures in the positive and negative directions. Greater opacity reflects a higher likelihood failure.

fall more directly in the negative direction, while some failures in the positive direction hesitate before falling. This may indicate that this particular controller learned by PPO is more robust to disturbances in the positive- θ direction than in the negative- θ direction.

Autonomous Vehicle Scenario: Sampled pedestrian trajectories that result in a collision are shown in Fig. 5. The HMC samples are subsampled to 500 trajectories for clarity. All trajectories end at the time of collision. HMC finds a wide variety of failure modes in terms of the trajectory of the pedestrian while the baselines focus on a single mode.

The more likely failure modes involve little change in the acceleration of the pedestrian. These failure trajectories are concentrated towards the y -axis, and introduce sensor noise to cause the AV to incorrectly estimate the relative distance or velocity to the pedestrian. Other failure modes involve accelerations of the pedestrian in the negative- x direction. Baseline methods were unable to sample from these modes. In extreme cases, the AV comes to a complete stop in front of the crosswalk and the pedestrian walks into the vehicle. This failure mode has been observed in previous work (Koren et al., 2018; Corso et al., 2019). While this pedestrian-induced failure might not be crucial to the safety of the AV policy, it highlights important questions for system designers such as accident blame that are not apparent from the baseline results.




Figure 5: Pedestrian trajectories resulting in collision for the autonomous vehicle scenario with HMC and baseline methods. HMC samples failures with a mixture of high sensor noise and lateral motion. Baseline methods only sample from the sensor noise failure mode. Greater opacity reflects a higher likelihood failure.




Figure 6: Failure trajectories for the lunar lander in altitude y and vertical velocity \dot{y} with HMC and baseline methods. Greater opacity reflects a higher likelihood failure.

Partially Observable Lunar Lander: Failure trajectories for the lunar lander are shown in Fig. 6. In the most likely failure mode, there is relatively small noise during the initial segment of the descent. A small amount of noise is added to the *AGL* measurement in the negative direction, making the lander believe that it is slightly closer to the ground than it really is. Next, a larger amount of noise is added to *AGL* in the positive direction, preventing the lander from increasing thrust in a critical moment. This highlights the ability of the proposed approach to find failures in non-Markov systems. The model-based approach is able to find diverse-likely failures by computing gradients through sequences of observations to find belief states that lead to failure.

6. Conclusion

Estimating the distribution over failures is key to the safe and confident deployment of autonomous systems in safety-critical domains. This work frames the estimation of the distribution over failure trajectories in sequential systems as Bayesian inference. We take a model-based approach, where system gradients are used in HMC to sample disturbance trajectories that lead to failure. Experiments against baseline approaches demonstrate that the proposed approach outperforms baselines in efficiently finding diverse, likely failure modes. The current approach is limited to systems compatible with AD, and is best suited when gradients are relatively inexpensive to compute. Future work will investigate using surrogate models (Qin et al., 2022) or stochastic gradient approximations (Chen et al., 2014) for systems not compatible with AD, and approximate methods like variational inference (Kucukelbir et al., 2017) when gradients are expensive.

Acknowledgments

This material is based upon work supported in part by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-2146755. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. This research was also supported in part by funding from Motional, Inc. This paper solely reflects the opinions and conclusions of its authors and not Motional or any other Motional entity.

References

- Takumi Akazaki, Shuang Liu, Yoriyuki Yamagata, Yihai Duan, and Jianye Hao. Falsification of cyber-physical systems using deep reinforcement learning. In *International Symposium on Formal Methods (FM)*, pages 456–465. Springer International Publishing, 2018.
- Rajeev Alur. *Principles of Cyber-Physical Systems*. MIT Press, 2015.
- Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3): 269–342, 2010.
- Claudine Badue, Rânik Guidolini, Raphael Vivacqua Carneiro, Pedro Azevedo, Vinicius B Cardoso, Avelino Forechi, Luan Jesus, Rodrigo Berriel, Thiago M Paixao, Filipe Mutz, et al. Self-driving cars: A survey. *Expert Systems with Applications*, 165:113816, 2021.
- Zdravko I Botev and Dirk 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.
- Tianqi Chen, Emily Fox, and Carlos Guestrin. Stochastic gradient Hamiltonian Monte Carlo. In *International Conference on Machine Learning (ICML)*, pages 1683–1691, 2014.
- Edmund M Clarke, Thomas A Henzinger, Helmut Veith, Roderick Bloem, et al. *Handbook of model checking*, volume 10. Springer, 2018.
- Anthony Corso, Peter Du, Katherine Driggs-Campbell, and Mykel J Kochenderfer. Adaptive stress testing with reward augmentation for autonomous vehicle validation. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 163–168. IEEE, 2019.
- Anthony Corso, Ritchie Lee, and Mykel J. Kochenderfer. Scalable autonomous vehicle safety validation through dynamic programming and scene decomposition. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6, 2020.
- Anthony Corso, Robert J Moss, Mark Koren, Ritchie Lee, and Mykel J Kochenderfer. A survey of algorithms for black-box safety validation of cyber-physical systems. *Journal of Artificial Intelligence Research*, 72(2005.02979):377–428, 2021.
- Jyotirmoy Deshmukh, Marko Horvat, Xiaoqing Jin, Rupak Majumdar, and Vinayak S Prabhu. Testing cyber-physical systems through Bayesian optimization. *ACM Transactions on Embedded Computing Systems (TECS)*, 16(5s):1–18, 2017.

- Tommaso Dreossi, Thao Dang, Alexandre Donz , James Kapinski, Xiaoqing Jin, and Jyotirmoy V Deshmukh. Efficient guiding strategies for testing of temporal properties of hybrid systems. In *NASA Formal Methods Symposium (NFM)*, pages 127–142, 2015.
- Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.
- Joel M Esposito, Jongwoo Kim, and Vijay Kumar. Adaptive RRTs for validating hybrid robotic control systems. In *Algorithmic Foundations of Robotics VI*, pages 107–121. Springer, 2004.
- Hong Ge, Kai Xu, and Zoubin Ghahramani. Turing: A language for flexible probabilistic inference. In *International Conference on Artificial Intelligence and Statistics*, pages 1682–1690. PMLR, 2018.
- Sebastian Geyer, Iason Papaioannou, and Daniel Straub. Cross entropy-based importance sampling using Gaussian densities revisited. *Structural Safety*, 76:15–27, 2019.
- Mark Girolami and Ben Calderhead. Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society. Series B, Statistical Methodology*, pages 123–214, 2011.
- Andrew D. Gordon, Thomas A. Henzinger, Aditya V. Nori, and Sriram K. Rajamani. Probabilistic programming. In *International Conference on Software Engineering (ICSE)*, page 167–181, 2014.
- Matthew D Hoffman, Andrew Gelman, et al. The no-u-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(1):1593–1623, 2014.
- Mike Innes, Alan Edelman, Keno Fischer, Chris Rackauckas, Elliot Saba, Viral B Shah, and Will Tebbutt. A differentiable programming system to bridge machine learning and scientific computing. *arXiv preprint arXiv:1907.07587*, 2019.
- Joost-Pieter Katoen. The probabilistic model checking landscape. In *ACM/IEEE Symposium on Logic in Computer Science*, pages 31–45, 2016.
- Youngjun Kim and Mykel J Kochenderfer. Improving aircraft collision risk estimation using the cross-entropy method. *Journal of Air Transportation*, 24(2):55–62, 2016.
- Mark Koren, Saud Alsaif, Ritchie Lee, and Mykel J Kochenderfer. Adaptive stress testing for autonomous vehicles. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 1–7. IEEE, 2018.
- Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M Blei. Automatic differentiation variational inference. *Journal of Machine Learning Research*, 18:1–45, 2017.
- Ritchie Lee, Mykel J Kochenderfer, Ole J Mengshoel, Guillaume P Brat, and Michael P Owen. Adaptive stress testing of airborne collision avoidance systems. In *Digital Avionics Systems Conference (DASC)*, pages 6C2–1. IEEE, 2015.
- Ritchie Lee, Ole J Mengshoel, Anshu Saksena, Ryan W Gardner, Daniel Genin, Joshua Silbermann, Michael Owen, and Mykel J Kochenderfer. Adaptive stress testing: Finding likely failure events with reinforcement learning. *Journal of Artificial Intelligence Research*, 69:1165–1201, 2020.

- Radford Neal. MCMC using Hamiltonian dynamics. In Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors, *Handbook of Markov Chain Monte Carlo*, chapter 5, pages 113–117. Chapman & Hall/CRC, 2011.
- Justin Norden, Matthew O’Kelly, and Aman Sinha. Efficient black-box assessment of autonomous vehicle safety. *arXiv preprint arXiv:1912.03618*, 2019.
- Matthew O’Kelly, Aman Sinha, Hongseok Namkoong, Russ Tedrake, and John C Duchi. Scalable end-to-end autonomous vehicle testing via rare-event simulation. *Advances in Neural Information Processing Systems (NIPS)*, 31, 2018.
- Xin Qin, Yuan Xian, Aditya Zutshi, Chuchu Fan, and Jyotirmoy V Deshmukh. Statistical verification of cyber-physical systems using surrogate models and conformal inference. In *International Conference on Cyber-Physical Systems (ICCPS)*, pages 116–126. IEEE, 2022.
- Daniel Rudoy and Patrick J Wolfe. Monte Carlo methods for multi-modal distributions. In *Asilomar Conference on Signals, Systems and Computers*, pages 2019–2023. IEEE, 2006.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Aman Sinha, Matthew O’Kelly, Russ Tedrake, and John C Duchi. Neural bridge sampling for evaluating safety-critical autonomous systems. In *Advances in Neural Information Processing Systems (NIPS)*, volume 33, pages 6402–6416, 2020.
- Anna Straubinger, Raoul Rothfeld, Michael Shamiyeh, Kai-Daniel Büchter, Jochen Kaiser, and Kay Olaf Plötner. An overview of current research and developments in urban air mobility—setting the scene for UAM introduction. *Journal of Air Transport Management*, 87:101852, 2020.
- Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62:1805, 2000.
- Cumhur Erkan Tuncali and Georgios Fainekos. Rapidly-exploring random trees for testing automated vehicles. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 661–666. IEEE, 2019.
- Ding Zhao, Xianan Huang, Huei Peng, Henry Lam, and David J LeBlanc. Accelerated evaluation of automated vehicles in car-following maneuvers. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):733–744, 2017.