

Enhanced Importance Sampling Through Latent Space Exploration in Normalizing Flows

Liam Anthony Kruse, Alexandros Tzikas, Harrison Delecki, Mansur Arief, Mykel Kochenderfer

Stanford University
Department of Aeronautics and Astronautics
Stanford, CA 94305 USA
{lkruse, alextzik, hdelecki, ariefm, mykel}@stanford.edu

Abstract

Importance sampling is a rare event simulation technique used in Monte Carlo simulations to bias the sampling distribution towards the rare event of interest. By assigning appropriate weights to sampled points, importance sampling allows for more efficient estimation of rare events or tails of distributions. However, importance sampling can fail when the proposal distribution does not effectively cover the target distribution. In this work, we propose a method for more efficient sampling by updating the proposal distribution in the latent space of a normalizing flow. Normalizing flows learn an invertible mapping from a target distribution to a simpler latent distribution. The latent space can be more easily explored during the search for a proposal distribution, and samples from the proposal distribution are recovered in the space of the target distribution via the invertible mapping. We empirically validate our methodology on simulated robotics applications such as autonomous racing and aircraft ground collision avoidance.

Introduction

Safety-critical applications such as autonomous driving or aircraft controller design heavily rely on simulations to enhance safety through testing in controlled environments. Potential failures can be identified through simulation and then addressed before real-world deployment, reducing the risk of accidents (Corso et al. 2021). Failures are often rare and safety thresholds are strict, so the events of interest—such as collisions or leaving a safe dynamic envelope—might be rarely encountered in simulation. *Importance sampling* (IS) is a variance reduction technique used in Monte Carlo simulations to bias the sampling distribution towards the rare event of interest (Owen 2013; Corso et al. 2021). IS uses a *proposal distribution* that focuses computational resources on scenarios likely to yield failure events, thus improving efficiency in failure detection. By assigning appropriate weights to sampled points, IS allows for more efficient estimation of the probability of failure compared to direct sampling from the target distribution.

Importance sampling can fail when the proposal distribution does not effectively cover the target distribution. If




Figure 1: Importance sampling in target space (top row) versus importance sampling in a flow’s latent space (bottom row). The target space proposal distribution generates many samples in the low-probability valley between the two failure modes, while the latent proposal generates samples that more closely map to the two failure regions.

the proposal distribution poorly matches the target distribution, the importance weights assigned to sampled points may become highly variable or imbalanced, leading to inaccurate failure estimates and increased estimate variance (Owen 2013). Furthermore, IS methods can miss less likely failure outcomes when the failure domain is multimodal, resulting in a biased estimate of the failure likelihood (Geyer, Papaioannou, and Straub 2019).

In this work we propose a technique to improve the efficiency of estimating the probability of failure by conducting importance sampling in the latent space of a normalizing flow. Normalizing flows are powerful generative models that learn an invertible mapping between a complex target density of interest and a simple, easy-to-evaluate latent density. We show that the distribution over simulation outcomes is mapped to a latent distribution that closely matches the IS prior, resulting in a balanced initial exploration over the space of possible outcomes. Moving forward, we refer to the space of simulation outcomes as the *target space*. We posit that the latent space of a normalizing flow is more ef-

fectively explored by IS methods than the target space. Furthermore, we observe that latent space IS provides empirical advantages over target space IS, such as improved sample efficiency and better coverage of failure events, as shown in Fig. 1. Our specific contributions include the following:

- We use the invertibility of normalizing flows to search for failure events in the latent space of flow models, demonstrating that importance sampling in latent space improves performance.
- We propose and justify an intuitive, easy-to-evaluate limit function formulation based on Löwner–John ellipsoids (Boyd and Vandenberghe 2004) to facilitate latent space exploration.
- We evaluate our approach on simulated robotics applications including autonomous racing and aircraft ground collision avoidance.

Related Work

A rich body of literature exists on importance sampling methods, reflecting their widespread adoption for applications such as structural reliability analysis (Kurtz and Song 2013; Papaioannou, Papadimitriou, and Straub 2016; Geyer, Papaioannou, and Straub 2019) and safety validation for autonomous vehicles (Huang et al. 2019; Corso et al. 2021). IS methods can fail when the target distribution is not adequately explored, resulting in an *inefficient* sampler, i.e., the effective sample size is small compared to the actual number of samples drawn (Corso et al. 2021). Our proposed technique enhances IS coverage by fitting proposal distributions in the latent space of a normalizing flow, using the fact that the latent density is usually simpler than the target density.

Normalizing flows have been used to provide expressive proposal distributions in IS methods (Müller et al. 2019; Gabrié, Rotskoff, and Vanden-Eijnden 2022; Samsonov et al. 2022). However, these approaches must interleave the sampling and training processes, resulting in complex training schemes and more strict assumptions to prove convergence (Samsonov et al. 2022). Our proposed methodology involves importance sampling in the latent space of a *pre-trained* flow model, providing more flexibility when the failure criteria or rare events of interest are updated yet the target distribution remains the same. Furthermore, our approach is beneficial if the flow is costly to train from scratch.

Researchers have recently identified the utility of implementing Monte Carlo methods in the favorable geometry of normalizing flow latent spaces. Hoffman et al. (2019) propose a Hamiltonian Monte Carlo algorithm for sampling in the latent space of an inverse autoregressive flow (Kingma et al. 2016), which can improve mixing speed. Coeurdoux, Dobeon, and Chainais (2023) propose a technique based on Langevin diffusion to correct for the topological mismatch between a latent unimodal distribution and a target distribution with disconnected support. Their goal is to limit out-of-distribution flow samples, whereas ours is to perform importance sampling for safety analyses. Noé et al. (2019) perform Metropolis Monte Carlo in the latent space of a Boltzmann generator to generate independent samples

of condensed matter systems and protein molecules. Meanwhile, Choi, Liao, and Ermon (2021) map two densities to a shared latent feature space to obtain more accurate density ratio estimates. Estimating density ratios is also a concern of Sinha et al. (2020), who perform bridge sampling after transforming the target space with a flow model.

Although importance sampling for rare events is the primary focus of this work, our methodology also lends itself to *conditional flow sampling*, wherein the mapping between the latent and target spaces is conditioned on an input (Winkler et al. 2019). Researchers have explored how the latent space can be partitioned to map different components of the input into disjoint regions in the target space (Dinh et al. 2019; Winkler et al. 2019). Whang, Lindgren, and Dimakis (2021) perform approximate conditional inference for image inpainting by composing two flow models, while Cannella, Soltani, and Tarokh (2020) define a Markov chain within a flow’s latent space to perform conditional image completion. Our approach can be viewed as an approximate conditional sampling scheme, enabling the flow to more efficiently generate samples that satisfy user-defined requirements.

Importance Sampling and Normalizing Flows

We outline the fundamental theory behind importance sampling and normalizing flows before justifying the decision to perform IS in the latent space of a pre-trained flow model.

Importance Sampling

Simulations allow engineers to assess the performance of algorithms and models in diverse scenarios, including rare or dangerous scenarios that would be costly to replicate in real-world testing. Assessing the probability of failure events can require a prohibitively large number of Monte Carlo simulations, especially if the event of interest is rare.

Consider an outcome space $\mathbf{x} \in \mathbb{R}^n$ with probability density function $p(\mathbf{x})$ and a *cost function* $f(\mathbf{x})$ such that a *failure event* occurs if and only if $f(\mathbf{x}) \leq 0$. The probability of failure P_F is given by the integral

$$P_F = \mathbb{E}_{p(\mathbf{x})} [\mathbb{1}\{f(\mathbf{x}) \leq 0\}] = \int \mathbb{1}\{f(\mathbf{x}) \leq 0\} \cdot p(\mathbf{x}) d\mathbf{x} \quad (1)$$

We can estimate P_F via Monte Carlo simulations by drawing N_s samples $\{\mathbf{x}_1, \dots, \mathbf{x}_{N_s}\}$ from $p(\mathbf{x})$ and taking the mean:

$$\hat{P}_F = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbb{1}\{f(\mathbf{x}_i) \leq 0\}. \quad (2)$$

This estimate is unbiased and has a coefficient of variation

$$\delta_{\hat{P}_F} = \sqrt{\frac{1 - P_F}{N_s P_F}}. \quad (3)$$

Since the coefficient of variation is inversely proportional to the failure probability, many samples might be required to come up with a precise estimate of P_F , especially if P_F is small (Papaioannou, Papadimitriou, and Straub 2016).

Importance sampling is a Monte Carlo simulation technique that aims to reduce the variance of \hat{P}_F by sampling

from an alternative sampling distribution—or *proposal distribution*—denoted by $q(\mathbf{x})$. So long as the support of $q(\mathbf{x})$ contains the failure domain, the probability of failure integral in Eq. (1) can be rewritten as

$$P_F = \int \frac{\mathbb{1}\{f(\mathbf{x}) \leq 0\} \cdot p(\mathbf{x})}{q(\mathbf{x})} \cdot q(\mathbf{x}) d\mathbf{x}. \quad (4)$$

The importance sampling estimate of P_F is given by

$$\hat{P}_F = \frac{1}{N_s} \sum_{i=1}^{N_s} \mathbb{1}\{f(\mathbf{x}_i) \leq 0\} \cdot \frac{p(\mathbf{x}_i)}{q(\mathbf{x}_i)} \quad (5)$$

where the samples are distributed according to the proposal distribution $q(\mathbf{x})$. Thus, an appropriate choice of proposal distribution can reduce the variance of the estimate of P_F .

Normalizing Flows

Normalizing flows (Rezende and Mohamed 2015) are a class of generative model used for density estimation and generative sampling. Normalizing flows transform a real vector \mathbf{u} sampled from an easy-to-evaluate *base distribution*, denoted by $p_u(\mathbf{u})$, through a transformation $\mathbf{z} = T(\mathbf{u})$ to produce a more expressive target density $p^*(\mathbf{z})$. A common choice of base distribution is a standard normal distribution. The transformation must be *invertible* and *differentiable*, i.e., a *diffeomorphism*. Imposing such topological constraints on a flow architecture ensures that the target density can be evaluated using the change of variables formula:

$$p^*(\mathbf{z}) = p_u(\mathbf{u}) |\det J_T(\mathbf{u})|^{-1} \quad \text{where } \mathbf{z} = T(\mathbf{u}). \quad (6)$$

The Jacobian of transformation T is denoted by J_T ; its determinant is a volume-correcting term that adjusts the probability density function of the transformed variable. The transformation (“flow”) itself is typically *composed* of D simpler transformations: $T = T_D \circ T_{D-1} \circ \dots \circ T_1$. Since the transformations are composable and each step (“layer”) is a diffeomorphism, we can set $\mathbf{z}_0 = \mathbf{u}$, $\mathbf{z}_d = T_d \circ \dots \circ T_1(\mathbf{z}_0)$ and compute the Jacobian-determinant in the log domain as

$$\log |\det J_T(\mathbf{u})| = \sum_{d=1}^D \log |\det J_{T_d}(\mathbf{z}_{d-1})|. \quad (7)$$

Figure 2 shows the two operations provided by a normalizing flow. The *forward* transformation is used when sampling, as an initial sample \mathbf{u} is drawn from the base distribution and cascaded through the composed transformations. The *inverse* transformation is used to perform density evaluation. The probability density function can be evaluated for an arbitrarily complex target distribution by iteratively computing the change of variables formula. Design considerations for both forward and inverse transformations are provided by Papamakarios et al. (2021).

Justification of Latent Space Sampling

Since normalizing flows learn an invertible mapping, we can fit proposal distributions over either the target space or the learned latent representation. As shown in Fig. 1, the target distribution can be highly non-isotropic in nature, with regions of extremely low density bordered by regions of very




Figure 2: A normalizing flow transforms a base distribution to a target distribution.

high likelihood. Thus, naively fitting a parametric proposal density can result in many generated samples that fall in regions of near-zero likelihood. Meanwhile, the latent distribution of normalizing flows is typically a standard normal, which is also a common choice for importance sampling priors (Papaoiannou, Papadimitriou, and Straub 2016). Probability mass in a small neighborhood around a latent data point is nearly isotropic, which facilitates exploration by isotropic proposals (Hoffman et al. 2019; Cannella, Soltani, and Tarokh 2020).

Methodology

In this section, we formalize our proposed technique for enhanced importance sampling and present our flow architecture, cost function formulation, and IS proposal-fitting methods.

Normalizing Flow Architecture

Recall that the flow transformation must be both invertible and differentiable. A sufficient condition for invertibility is enforcing the transformation to be monotonic. Durkan et al. (2019) propose the use of monotonic rational quadratic splines as building blocks for the transform. In this work we use piecewise rational quadratic *coupling* transforms to compose the flow models. Coupling transforms operate by splitting the input into two parts and applying an invertible function to one part while leaving the other part unchanged, thus ensuring efficient inversion and sampling capabilities (Dinh, Sohl-Dickstein, and Bengio 2016; Papamakarios et al. 2021). The flow parameters are optimized by minimizing the forward KL divergence between the learned and target distributions.

Cost Function Formulation

The cost function should be continuous and bias the search over simulator outcomes towards the rare or failure events of interest (Corso et al. 2021). However, computing the distance function for an arbitrary set is a hard problem, and even determining if an outcome falls inside or outside a set can be computationally expensive (Hormann and Agathos 2001). We propose a cost function geometry based

on Löwner–John ellipsoids, i.e., the minimum volume ellipsoid that contains a set of points \mathcal{X} (Boyd and Vandenberghe 2004). We posit that minimum volume ellipsoids are a reasonable choice of cost function since 1) they overapproximate the convex hull of \mathcal{X} , which is desirable for enforcing conservative safety thresholds and 2) Gaussian mixture models are commonly used for proposal distributions, and each mixture component has ellipsoidal level sets.

Consider a finite set of points $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\} \subseteq \mathbb{R}^n$. The problem of finding the minimum volume ellipsoid that covers \mathcal{X} is a convex optimization problem (Boyd and Vandenberghe 2004):

$$\begin{aligned} & \text{minimize} && \log \det \mathbf{A}^{-1} \\ & \text{subject to} && \|\mathbf{A}\mathbf{x}_i + \mathbf{b}\|_2 \leq 1, \quad i = 1, \dots, m \end{aligned} \quad (8)$$

where the variables are a symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$. This formulation is attractive because the Löwner–John ellipsoid can be found quickly using standard optimization solvers such as `CVXPY` (Diamond and Boyd 2016).

A question that naturally arises is how to obtain a set of points \mathcal{X} that is representative of the failure region. If the data used to train the flow model is available, then one solution is to construct \mathcal{X} from the failures in this dataset. However, if the dataset is relatively small or the failure events are extremely rare, then the set \mathcal{X} might not adequately cover the entire failure domain. In this event, the overapproximating nature of Löwner–John ellipsoids is advantageous. Another solution is to specify the set \mathcal{X} manually based on user insight (e.g., defining the corners of a safety threshold hypercube).

Once \mathcal{X} is obtained, the representative points are mapped deterministically to a set \mathcal{U} in latent space via the invertible flow mapping. Solving Eq. (8) with input \mathcal{U} yields a Löwner–John ellipsoid in latent space. The Mahalanobis distance d_M between the ellipsoidal region and a sample from the IS proposal density is easily computed. If $d_M \leq 0$, the sample lies within the ellipsoid and is classified as a failure event.

Importance Sampling Methods

In this work we evaluate two importance sampling methods; however, any IS algorithm could be used in practice (Owen 2013; Corso et al. 2021). The first algorithm is the *cross entropy* (CE) method (De Boer et al. 2005; Geyer, Papaioannou, and Straub 2019), which attempts to learn the parameters of a parametric proposal distribution that minimizes the KL divergence between the optimal IS density and the proposal. The CE method introduces a series of intermediate failure domains that gradually approach the true failure domain. At each step, the intermediate failure region is defined such that $\rho \cdot N_s$ samples fall in the region, where the ρ -quantile is chosen by the user. The proposal distribution parameters are then fit via maximum likelihood estimation over these samples. The expectation-maximization algorithm is often used to fit the search distribution, though it must be adjusted to account for importance-weighted samples (Geyer, Papaioannou, and Straub 2019).

The second method is *sequential importance sampling* (SIS) (Del Moral, Doucet, and Jasra 2006; Papaioannou, Pa-

Algorithm 1: Latent Space Importance Sampling

Require: flow T , failure set $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$, proposal $q_0(\mathbf{u}; \theta_0)$, latent flow distribution $p(\mathbf{u}) = \mathcal{N}(\mathbf{u}; \mathbf{0}, \mathbf{I})$
Ensure: probability of failure P_F

- 1: **procedure** LATENT IMPORTANCE SAMPLING
- 2: $\mathcal{U} \leftarrow T^{-1}(\mathcal{X})$
- 3: $\mathbf{A}, \mathbf{b} \leftarrow$ solve for Löwner–John ellipsoid of \mathcal{U}
- 4: $\Sigma = (\mathbf{A}^\top \mathbf{A})^{-1}$, $\boldsymbol{\mu} = -(\Sigma \mathbf{A}^\top) \mathbf{b}$
- 5: define function $f(\mathbf{u}) = \sqrt{(\mathbf{u} - \boldsymbol{\mu}) \Sigma^{-1} (\mathbf{u} - \boldsymbol{\mu})^\top}$
- 6: **for** $k \leftarrow 1 : k_{\max}$
- 7: elite samples \mathbf{u}_e , importance weights $\mathbf{w} \leftarrow$
- 8: ISMETHOD($f(\mathbf{u}), q_k(\mathbf{u}; \theta_k), p(\mathbf{u})$)
- 9: **while** not converged
- 10: $\theta_{k+1} \leftarrow$ fit proposal parameters with \mathbf{u}_e, \mathbf{w}
- 11: compute P_F with Eq. (5)

padimitriou, and Straub 2016). Like the CE method, SIS introduces a series of intermediate failure distributions that gradually approach the optimal IS density. Samples for each intermediate distribution are obtained by resampling weighted particles from the previous distribution and then moved to regions of high likelihood under the next failure distribution via Markov Chain Monte Carlo. In this work, we use a conditional sampling Metropolis–Hastings algorithm to move the samples (Papaioannou, Papadimitriou, and Straub 2016).

Algorithm

Algorithm 1 presents the proposed latent IS methodology for estimating P_F . The algorithm takes as input a flow model and an importance sampling algorithm and outputs an estimate of P_F . Furthermore, failure events can be easily generated after learning the proposal distribution by generating samples in latent space and mapping them back to target space. Without loss of generality, the user can input a set of failure sets $\{\mathcal{X}_i\}_{i=0}^n$, $\mathcal{X}_i = \{\mathbf{x}_1^i, \dots, \mathbf{x}_m^i\}$, with each set corresponding to a failure mode. The Löwner–John ellipsoid is solved for each failure set, and the cost function computes the Mahalanobis distance for each ellipsoid and returns the minimum value.

Experiments

This section presents our simulated robotics datasets and evaluation metrics before discussing experimental results.

Data Simulators

We use three autonomous system simulators to validate our proposed approach.

Nonholonomic Robot Consider a nonholonomic robot that moves in two dimensions. The three-dimensional robot state s is defined as

$$s = [x \quad y \quad \theta]^\top,$$

where x and y are the x- and y-positions of the vehicle and θ is the heading angle. The control vector a is

$$a = [v \quad \alpha]^\top$$

with velocity v and angular rate α in radians per second. We simulate noisy dynamics over a $T = 40$ second horizon, holding the inputs at constant magnitude for the duration of each simulation. However, to induce a multimodal outcome, we stochastically flip the sign of the angular rate input 15 seconds into each trial. We train a flow on 10^5 trials and save 10^6 simulations for Monte Carlo evaluations.

Cornering Racecar We next simulate data from a nonlinear single-track racecar modeled with a variant of the Fiala brush tire model (Subosits and Gerdes 2021). The car uses model predictive path integral control (Williams et al. 2018) to drift around a corner without spinning out of control. The eight-dimensional car state s is defined as

$$s = [x \ y \ \psi \ v_x \ v_y \ \dot{\psi} \ \delta \ F_p]^\top,$$

where x and y are the x- and y-positions, v_x and v_y are the longitudinal and lateral velocities, ψ and $\dot{\psi}$ are the yaw and yaw rate, δ is the steering angle, and F_p is an input from the pedals. We use the experimental setup of Asmar et al. (2023), who provide an open-source Julia repository for their experiments.¹ The flow is trained on 5×10^5 trials, with 1.05×10^6 simulations saved for Monte Carlo evaluations.

F-16 Ground Collision Avoidance Finally, we simulate an F-16 fighter jet controlled by a ground collision avoidance system (GCAS). We use the dynamics model introduced by Heidlauf et al. (2018) and the JAX code implementation² by So and Fan (2023). The F-16 begins in a dive towards the ground; the GCAS system then rolls the aircraft until the wings are level and pulls the nose above the horizon. We add noise to the F-16’s initial state, thus simulating an envelope of possible start positions from which the GCAS must execute a recovery. The post-processed state vector is 12-dimensional, including the aircraft’s roll, pitch, yaw, altitude, and airspeed. For a full list and definition of state variables, please refer to Heidlauf et al. (2018). Note that we remove the engine power lag variable (which is discrete), the stability roll rate and side acceleration/yaw rate integrators (which are observed to be close to zero), and the angle of attack variable. We train on 10^6 trials and save 10^7 simulations for Monte Carlo evaluations.

Metrics

We compute a reference failure probability for each dataset using the Monte Carlo estimate given in Eq. (2) and then calculate the importance sampling estimate \hat{P}_F using Eq. (5). We return the relative error between these two values; a positive value indicates that the estimated failure likelihood is an overestimate of the true value. Next, we compute a series of metrics to evaluate the quality of the learned proposal density:

- *Density*: The density metric proposed by Naeem et al. (2020) rewards the proposal distribution for generating

¹<https://github.com/sisl/MPOPIS>

²<https://github.com/MIT-REALM/jax-f16>

samples in regions where real data points (i.e., true failure events) are closely packed. A *higher* value indicates better performance.

- *Coverage*: The coverage metric proposed by Naeem et al. (2020) measures the fraction of real samples whose neighborhood contains at least one generated sample; it is useful for detecting mode dropping. A *higher* value indicates better coverage.
- *Average negative log likelihood (NLL)*: An evaluation batch of size N_{eval} is drawn from the learned proposal distribution and the negative log-likelihood of each data point is computed according to Eq. (6) and Eq. (7). A *lower* value indicates that the samples more closely match the learned flow density.

Lastly, we report the average number of samples until convergence, \bar{N}_{total} , to evaluate sample efficiency. A *lower* value indicates that the method requires fewer samples and function evaluations to find a sufficient proposal distribution.

Experimental Setup

We perform cross-entropy importance sampling and sequential importance sampling on all three datasets; furthermore, we evaluate each method in both latent space and target space. For the target space experiments, we compute the Löwner–John ellipsoids on $\{\mathcal{X}_i, i = 0, \dots, n\}$ directly. In the latent experiments, we first compute $\mathcal{U} = T^{-1}(\mathcal{X})$ and then solve Eq. (8) with constraints based on the points in \mathcal{U} .

For the nonholonomic robot data, we construct \mathcal{X}_1 as a unit cube such that $x \in [-1.0, -2.0]$, $y \in [-2.25, -3.25]$, $\theta \in [1.25, 2.25]$, and \mathcal{X}_2 as a unit cube such that $x \in [0.75, 1.75]$, $y \in [-3.25, -4.25]$, $\theta \in [-1.0, -2.0]$. Figure 3 visualizes these regions in target space (left) and latent space (right).

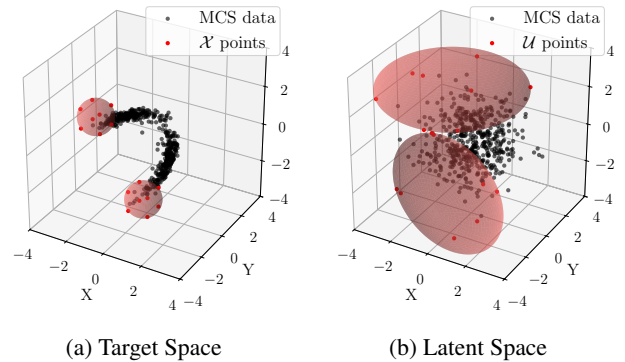


Figure 3: Failure regions for the nonholonomic robot shown in target space and latent space. The points in \mathcal{X}_i are mapped to latent space and the Löwner–John ellipsoids are re-computed.

For the racecar experiments, a small subset of the training data is considered to build the representative failure regions. \mathcal{X}_1 is constructed from the datapoints with $x > 0.0$, $\psi > 2.75$ (samples with a large yaw angle) and \mathcal{X}_2 is constructed from the datapoints with $x > 1.5$, $\dot{\psi} < -2.25$ (samples

| Method | $(\hat{P}_F - P_F)/P_F$ | avg NLL(\hat{x}) | coverage | density | \bar{N}_{total} |
|-------------------------------------------------------------------------|-------------------------|-----------------------|----------------------|----------------------|--------------------------|
| <i>Nonholonomic Robot (3D): $P_f = 0.0072$</i> | | | | | |
| Latent-CE | 0.462 ± 0.056 | 3.468 ± 0.299 | 0.787 ± 0.039 | 0.983 ± 0.027 | 6600.0 |
| Latent-SIS | 0.496 ± 0.210 | 3.107 ± 0.137 | 0.704 ± 0.033 | 0.990 ± 0.010 | 8520.0 |
| Target-CE | -0.924 ± 0.007 | 9.066 ± 1.122 | 0.486 ± 0.065 | 0.482 ± 0.078 | 9000.0 |
| Target-SIS | -0.913 ± 0.018 | 8.861 ± 0.475 | 0.512 ± 0.042 | 0.463 ± 0.034 | 10540.0 |
| <i>Cornering Racecar (8D): $P_f = 0.0062$</i> | | | | | |
| Latent-CE | -0.040 ± 0.012 | 8.469 ± 0.256 | 0.813 ± 0.010 | 1.057 ± 0.029 | 39900.0 |
| Latent-SIS | -0.054 ± 0.082 | 8.690 ± 0.117 | 0.789 ± 0.012 | 1.083 ± 0.017 | 51000.0 |
| Target-CE | -0.957 ± 0.003 | 17.206 ± 0.744 | 0.579 ± 0.040 | 0.542 ± 0.027 | 49900.0 |
| Target-SIS | -0.952 ± 0.008 | 16.246 ± 0.495 | 0.614 ± 0.046 | 0.605 ± 0.031 | 62900.0 |
| <i>F-16 Ground Collision Avoidance (12D): $P_f = 0.0043$</i> | | | | | |
| Latent-CE | -0.611 ± 0.001 | 32.443 ± 0.058 | 0.584 ± 0.007 | 0.663 ± 0.006 | 400000.0 |
| Latent-SIS | -0.664 ± 0.004 | 32.118 ± 0.079 | 0.571 ± 0.007 | 0.665 ± 0.007 | 505000.0 |
| Target-CE | -1.000 ± 0.000 | 38.136 ± 0.103 | 0.042 ± 0.003 | 0.014 ± 0.001 | 900000.0 |
| Target-SIS | -1.000 ± 0.000 | 37.024 ± 0.208 | 0.051 ± 0.003 | 0.016 ± 0.001 | 1205000.0 |

Table 1: Experimental Results

with a low yaw rate). Likewise, a small subset of the training data is used to build \mathcal{X}_1 and \mathcal{X}_2 for the F-16 experiments. \mathcal{X}_1 is obtained from the datapoints with a pitch $\theta > 1.45$ (high pitch) and \mathcal{X}_2 is constructed from the datapoints with an altitude < -2.45 (low altitude). Note that each state variable was normalized to have zero mean and unit variance. Code to reproduce the experimental results is available at <https://github.com/sisl/LatentImportanceSampling>.

Results

Table 1 presents the experimental results across the three datasets. We run 100 trials for the robot and racecar datasets and 20 trials for the F-16 dataset, recording the metric means and standard deviations. The latent sampling methods consistently achieve the lowest average negative log-likelihood values, indicating that the samples generated in latent space and pushed forward through the flow transformation more closely match the learned target density than points sampled directly in target space. Likewise, the latent-space IS methods achieve higher density and coverage scores than the target-space IS methods. The higher density scores indicate that latent IS places more samples in regions with a higher density of actual failure events. The higher coverage scores show that latent IS does a better job of finding all failure modes and generating samples in regions of the simulator outcome space that contain failure events. Performing the cross-entropy method in latent space results in the smallest average number of samples until convergence across all datasets, resulting in the lowest computational cost.

Discussion and Future Work

Importance sampling is a powerful method for computing the probability of rare events and validating autonomous systems (Corso et al. 2021). In this work we present a technique to improve importance sampling by first transforming the data with a normalizing flow. The invertible flow transformation warps non-isotropic target densities into a latent

density which is more isotropic and easier to explore. We also propose an intuitive cost-function formulation that is simple to evaluate, even after undergoing an arbitrarily complex transformation to latent space. We experimentally show that conducting IS in latent space results in failure samples that more closely match the true distribution of failure events. Samples generated in latent space can be easily recovered in target space via the deterministic and invertible flow mapping. Our latent IS methods outperform target IS methods on a range of simulated autonomous systems.

Future work will provide a theoretical analysis of the benefits of latent IS, investigating the impact of maximum likelihood training on the warped geometry of the latent failure regions. Alternative cost-function formulations will be explored; for example, maximum-volume inscribed ellipsoids could result in P_F estimates that are less conservative. Finally, the latent IS methods presented in this paper will be used to validate real-world, black-box autonomous systems with high-fidelity simulators.

Acknowledgements

Toyota Research Institute (TRI) provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity. The NASA University Leadership Initiative provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not any NASA entity.

References

- Asmar, D. M.; Senanayake, R.; Manuel, S.; and Kochenderfer, M. J. 2023. Model predictive optimized path integral strategies. In *IEEE International Conference on Robotics and Automation (ICRA)*, 3182–3188. IEEE.
- Boyd, S. P.; and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.

- Cannella, C.; Soltani, M.; and Tarokh, V. 2020. Projected Latent Markov Chain Monte Carlo: Conditional Sampling of Normalizing Flows. In *International Conference on Learning Representations (ICLR)*.
- Choi, K.; Liao, M.; and Ermon, S. 2021. Featurized density ratio estimation. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, 172–182. PMLR.
- Coeurdoux, F.; Dobigeon, N.; and Chainais, P. 2023. Normalizing flow sampling with Langevin dynamics in the latent space. *arXiv preprint arXiv:2305.12149*.
- Corso, A.; Moss, R.; Koren, M.; Lee, R.; and Kochenderfer, M. 2021. A survey of algorithms for black-box safety validation of cyber-physical systems. *Journal of Artificial Intelligence Research*, 72: 377–428.
- De Boer, P.-T.; Kroese, D. P.; Mannor, S.; and Rubinstein, R. Y. 2005. A tutorial on the cross-entropy method. *Annals of Operations Research*, 134: 19–67.
- Del Moral, P.; Doucet, A.; and Jasra, A. 2006. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3): 411–436.
- Diamond, S.; and Boyd, S. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83): 1–5.
- Dinh, L.; Sohl-Dickstein, J.; and Bengio, S. 2016. Density estimation using Real NVP. In *International Conference on Learning Representations (ICLR)*.
- Dinh, L.; Sohl-Dickstein, J.; Larochelle, H.; and Pascanu, R. 2019. A RAD approach to deep mixture models. *Deep Generative Models for Highly Structured Data Workshop, International Conference on Learning Representations (ICLR)*.
- Durkan, C.; Bekasov, A.; Murray, I.; and Papamakarios, G. 2019. Neural spline flows. *Advances in Neural Information Processing Systems (NeurIPS)*, 32.
- Gabrié, M.; Rotskoff, G. M.; and Vanden-Eijnden, E. 2022. Adaptive Monte Carlo augmented with normalizing flows. *Proceedings of the National Academy of Sciences*, 119(10): e2109420119.
- Geyer, S.; Papaioannou, I.; and Straub, D. 2019. Cross entropy-based importance sampling using Gaussian densities revisited. *Structural Safety*, 76: 15–27.
- Heidlauf, P.; Collins, A.; Bolender, M.; and Bak, S. 2018. Verification Challenges in F-16 Ground Collision Avoidance and Other Automated Maneuvers. In *International Workshop on Applied Verification for Continuous and Hybrid Systems (ARCH)*, 208–217.
- Hoffman, M.; Sountsov, P.; Dillon, J. V.; Langmore, I.; Tran, D.; and Vasudevan, S. 2019. Neutralizing bad geometry in Hamiltonian Monte Carlo using neural transport. *Symposium on Advances in Approximate Bayesian Inference*.
- Hormann, K.; and Agathos, A. 2001. The point in polygon problem for arbitrary polygons. *Computational Geometry*, 20(3): 131–144.
- Huang, Z.; Arief, M.; Lam, H.; and Zhao, D. 2019. Evaluation uncertainty in data-driven self-driving testing. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 1902–1907. IEEE.
- Kingma, D. P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; and Welling, M. 2016. Improved variational inference with inverse autoregressive flow. *Advances in Neural Information Processing Systems (NIPS)*, 29.
- Kurtz, N.; and Song, J. 2013. Cross-entropy-based adaptive importance sampling using Gaussian mixture. *Structural Safety*, 42: 35–44.
- Müller, T.; McWilliams, B.; Rousselle, F.; Gross, M.; and Novák, J. 2019. Neural importance sampling. *ACM Transactions on Graphics (ToG)*, 38(5): 1–19.
- Naeem, M. F.; Oh, S. J.; Uh, Y.; Choi, Y.; and Yoo, J. 2020. Reliable fidelity and diversity metrics for generative models. In *International Conference on Machine Learning (ICML)*, 7176–7185. PMLR.
- Noé, F.; Olsson, S.; Köhler, J.; and Wu, H. 2019. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457): eaaw1147.
- Owen, A. B. 2013. *Monte Carlo Theory, Methods and Examples*. Stanford University. <https://artowen.su.domains/mc/>.
- Papaioannou, I.; Papadimitriou, C.; and Straub, D. 2016. Sequential importance sampling for structural reliability analysis. *Structural Safety*, 62: 66–75.
- Papamakarios, G.; Nalisnick, E.; Rezende, D. J.; Mohamed, S.; and Lakshminarayanan, B. 2021. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57): 1–64.
- Rezende, D.; and Mohamed, S. 2015. Variational inference with normalizing flows. In *International Conference on Machine Learning (ICML)*, 1530–1538. PMLR.
- Samsonov, S.; Lagutin, E.; Gabrié, M.; Durmus, A.; Naumov, A.; and Moulines, E. 2022. Local-global MCMC kernels: the best of both worlds. *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 5178–5193.
- Sinha, A.; O’Kelly, M.; Tedrake, R.; and Duchi, J. C. 2020. Neural bridge sampling for evaluating safety-critical autonomous systems. *Advances in Neural Information Processing Systems (NeurIPS)*, 33: 6402–6416.
- So, O.; and Fan, C. 2023. Solving Stabilize-Avoid Optimal Control via Epigraph Form and Deep Reinforcement Learning.
- Subosits, J. K.; and Gerdes, J. C. 2021. Impacts of model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers. *IEEE Transactions on Intelligent Vehicles*, 6(3): 546–558.
- Wang, J.; Lindgren, E.; and Dimakis, A. 2021. Composing normalizing flows for inverse problems. In *International Conference on Machine Learning (ICML)*, 11158–11169. PMLR.
- Williams, G.; Drews, P.; Goldfain, B.; Reh, J. M.; and Theodorou, E. A. 2018. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6): 1603–1622.
- Winkler, C.; Worrall, D.; Hoogetboom, E.; and Welling, M. 2019. Learning Likelihoods with Conditional Normalizing Flows. *arXiv preprint arXiv:1912.00042*.