

# Physics-Informed Representation and Learning: Control and Risk Quantification

Zhuoyuan Wang<sup>1</sup>, Reece Keller<sup>1</sup>, Xiyu Deng<sup>1</sup>, Kenta Hoshino<sup>2</sup>, Takashi Tanaka<sup>3</sup>, Yorie Nakahira<sup>1</sup>

<sup>1</sup>Carnegie Mellon University

<sup>2</sup>Kyoto University

<sup>3</sup>University of Texas at Austin

{zhuoyuaw, rdkeller, xiyud}@andrew.cmu.edu, hoshino@i.kyoto-u.ac.jp, ttanaka@utexas.edu, yorie@cmu.edu

## Abstract

Optimal and safety-critical control are fundamental problems for stochastic systems, and are widely considered in real-world scenarios such as robotic manipulation and autonomous driving. In this paper, we consider the problem of efficiently finding optimal and safe control for high-dimensional systems. Specifically, we propose to use dimensionality reduction techniques from a comparison theorem for stochastic differential equations together with a generalizable physics-informed neural network to estimate the optimal value function and the safety probability of the system. The proposed framework results in substantial sample efficiency improvement compared to existing methods. We further develop an autoencoder-like neural network to automatically identify the low-dimensional features of the system to enhance the ease of design for system integration. We also provide experiments and quantitative analysis to validate the efficacy of the proposed method. Source code is available at <https://github.com/jacobwang925/path-integral-PINN>.

## Introduction

Optimal control and safety-critical control are the two central concerns for autonomous systems. These concerns are particularly pronounced in real-world applications, *e.g.*, manufacturing robots and autonomous cars. The operational environment often introduces stochastic noise, which compounds the difficulties of achieving optimal performance and ensuring safety. Traditional deterministic methods prove inadequate for stochastic dynamics (Katsoulakis and Vilanova 2020). Furthermore, many real-world systems are characterized by high-dimensional state spaces (*e.g.*, multi-agent systems), leading to substantial computational burdens when devising optimal and safe control strategies.

Previous work on stochastic optimal control deals with diverse uncertainty and randomness, but these methods are not efficient for high-dimensional systems because forward rollouts and backward dynamic programming requires computation that scales exponentially with the dimension of the state (Gorodetsky, Karaman, and Marzouk 2018; Frankowska and Zhang 2020). Previous studies on stochastic safe control aim to control the level of risk in the system and ensure the probability of safety does not decay

over time (Samuelson and Yang 2018; Gómez et al. 2016). These methods rely on accurate estimations of the probability of risk in the system, and standard methods for such risk estimation are computationally heavy, especially for high-dimensional systems. For both problems, existing methods require computation that scales exponentially or linearly with the time horizon of interest. To the best of our knowledge, there is no study that provides estimation of the value function or risk probability in long time horizons without introducing additional computation, which can be very beneficial for systems with long-term performance requirements.

In order to address these challenges in computation, this study proposes a unified framework to efficiently estimate the value function and safety probability of high-dimensional stochastic systems. The method leverages a comparison theorem (Yamada 1973) to find low-dimensional representations of the value function and safety probability, and uses such low-dimensional features to construct low-dimensional partial differential equations (PDEs) for the value function and safety probability calculation in order to reduce the dimension of the problem for efficient computation. We further propose a physics-informed neural network (PINN) to solve these PDEs for better sample complexity and generalization abilities. We also propose an autoencoder-like neural network to automatically identify low-dimensional features of the system. Fig. 1 shows the overall diagram of the proposed method. The advantages of the proposed method are

1. A unified framework for accurate estimation of value functions and safety probabilities of stochastic systems (Fig. 3a).
2. Efficient estimation with much lower sample complexity for high-dimensional systems (Fig. 5).
3. Generalization to unseen regions in the state space and longer time horizons (Fig. 4).
4. Intuitive plug and use with automatic feature identification (Fig. 3b).

## Related Work

### Path Integral Optimal Control

Path integral control generally refers to numerical methods to solve a stochastic optimal control problem by repeatedly performing forward Monte Carlo rollouts of open-loop dynamics. The original derivation of the path integral




Figure 1: The overall diagram of the proposed method for stochastic optimal value function and safety probability estimation.

control algorithm (Kappen 2005) relied on the Feynman-Kac lemma, whereas an alternative derivation without the Feynman-Kac lemma became available later (Theodorou and Todorov 2012) based on the variational approach. In both derivations, the path integral method is restricted to the class of stochastic optimal control problems whose Hamilton–Jacobi–Bellman (HJB) PDEs are linearizable, but generalizations have been considered in (Satoh, Kappen, and Saeki 2016; Williams et al. 2017). Various implementations such as path integral for policy improvements (Theodorou, Buchli, and Schaal 2010) and receding horizon implementations (Williams, Aldrich, and Theodorou 2017) have been widely used. Path integral control has also been applied to constrained systems and systems with non-differentiable dynamics (Satoh and Kappen 2020; Carius et al. 2022). Path integral control for risk minimization in robot navigation has been considered in (Patil et al. 2022), which requires solving PDEs whose dimension scales with the size of the system. Here, we show that the value function and safety probability can be bounded exactly by the solution of low-dimensional PDEs regardless of the system dimension.

### Risk Quantification for Safe Control

Risk quantification is the key enabler for many long-term stochastic safe control methods (Wang et al. 2021; Jing and Nakahira 2022). Existing methods often use rare event simulation through Monte Carlo (MC) and importance sampling to estimate the long-term risk in stochastic systems (Botev, L’Ecuyer, and Tuffin 2013; Hanna, Niekum, and Stone 2021; Stadie et al. 2018; Madhushani et al. 2021). The subset simulation calculates the risk probability conditioned on intermediate failure events for improved sample efficiency (Huang, Chen, and Zhu 2016; Zhao and Wang 2022; Rashki 2021). Probabilistic reachability estimates the risk of controllers in stochastic systems by propagating the estimated risk backwards over time (Hewing and Zeilinger 2018; Bansal et al. 2020; Huh and Yang 2020). MC techniques typically require samples to cover states and evaluate the risk over the time horizon. PDE techniques are also used to understand probabilistic values in stochastic systems (Chern et al. 2021; Feng 2014), but numerical PDE techniques such as finite difference, finite element, and finite volume methods are less scalable than MC methods. Probability bounds and martingale inequalities have been used to approximate risk probabilities for certain classes of systems (Clark 2019; Yaghoubi et al. 2020; Santoyo, Dutreix, and Coogan 2021; Cheng et al.

2020; Meng and Liu 2022; Nishimura and Hoshino 2023). The large deviation is another standard approach that can be adapted to the safe control area, which allows evaluating the probability of the state of a stochastic differential equation that exists from a given region (Bressloff and Newby 2014; Bertini, Faggionato, and Gabrielli 2015). Nonetheless, most existing methods suffer from the curse of dimensionality. In this work, we solve the high-dimensional safety probability estimation problem by finding effective low-dimensional representations, and we are able to reduce the sample complexity by orders compared to MC methods.

### Physics-informed Neural Networks

Physics-informed neural networks (PINNs) are neural networks that are trained to solve supervised learning tasks while respecting any given laws of physics described by general nonlinear PDEs (Raissi, Perdikaris, and Karniadakis 2019). PINNs take both data and the physics model of the system into account, and are able to solve the forward problem of getting PDE solutions, and the inverse problem of discovering underlying governing PDEs from data. PINNs have been widely used in power systems (Misyris, Venzke, and Chatzivasileiadis 2020), fluid mechanics, (Cai et al. 2022) medical care (Sahli Costabal et al. 2020), and risk quantification (Han, Jentzen, and Weinan 2018; Pereira et al. 2021; Wang and Nakahira 2023). We leverage the generalization ability of PINNs to efficiently estimate value functions and safety probabilities in unseen regions of the state space to further enhance the sample complexity of the proposed method.

## Problem Formulation

### System Dynamics

Consider the following class of nonlinear stochastic dynamical systems defined on a probability space  $(\Omega, F, P)$ :

$$dx_t = f(x_t) dt + \sigma(x_t)(u_t dt + dw_t) \quad (1)$$

Here,  $x_t \in \mathcal{X} \subset \mathbb{R}^n$ ,  $0 \leq t \leq T$  is the state of the system and  $u_t \in \mathcal{U} \subset \mathbb{R}^m$ ,  $0 \leq t \leq T$  is the control input,  $w_t$  is the  $m$ -dimensional standard Brownian motion in the probability measure  $P$  and  $\sigma$  is the diffusion coefficient. The role of the controller is to apply the control input  $u_t$  based on a state feedback policy *i.e.*,  $u_t$  is measurable with respect to the filtration  $F^{x_t}$  generated by  $\{x_\tau\}_{0 \leq \tau \leq t}$ . Suppose that  $Q$  is an alternative probability measure in which  $u_t \equiv 0$ . Correspondingly, we have

$$d\tilde{w}_t = u_t dt + dw_t, \tilde{w}_0 = 0 \quad (2)$$

is the standard Brownian motion.

### Stochastic Optimal Control

Consider the running cost defined as

$$w(x_t, u_t) = c(x_t) + \frac{1}{2} \|u_t\|^2, \quad (3)$$

where  $c : \mathbb{R}^n \rightarrow \mathbb{R}$ . The stochastic optimal control problem aims to find the optimal value function

$$V(x, t) := \min_u \mathbb{E}^P \left[ \int_t^T w(x_\tau, u_\tau) d\tau + c(x_T) \mid x_t = x \right], \quad (4)$$

which explicitly yields the optimal control as

$$u_t^* = -\sigma(x_t)^\top \nabla_x V(x_t, t). \quad (5)$$

A notable feature of the stochastic optimal control problems is the applicability of Monte Carlo-based numerical solution strategy, which we call the path-integral method (Thijssen and Kappen 2015). For each time  $t \in [0, T)$  and the state  $x \in \mathbb{R}^n$ , the path-integral method allows the control agent to compute the optimal input  $u_t^*$  by evaluating the path integrals along randomly generated state trajectories  $x_\tau, t \leq \tau \leq T$  starting from  $x_t = x$ . From existing results on KL control and free energy (Fleming and Soner 2006; Boué and Dupuis 1998; Theodorou and Todorov 2012), the value function (4) can be solved explicitly as

$$V(x, t) = -\log \mathbb{E}^Q \left[ \exp \left( -\int_t^T c(x_\tau) d\tau - c(x_T) \right) \mid x_t = x \right]. \quad (6)$$

Since the right hand side of (6) contains the expectation with respect to  $Q$ , one can consider approximating it by Monte Carlo simulation as

$$V(x, t) \approx -\log \left[ \frac{1}{N} \sum_{i=1}^N \exp \left( -\int_t^T c(x_\tau^i) d\tau - c(x_T^i) \right) \right], \quad (7)$$

where  $\{x_\tau^i, t \leq \tau \leq T\}_{i=1}^N$  are randomly drawn sample paths from distribution  $Q$ . Since  $\tilde{w}_t$  is the standard Brownian motion under  $Q$ , such sample paths can be obtained by simply simulating the uncontrolled system  $dx_t = f(x_t)dt + \sigma(x_t)d\tilde{w}_t$ . Note that  $Q$  is the uncontrolled process and is easy to sample from, while the value function given by (7) is associated with the optimal control. This is the path integral control method, and is widely adopted for low-dimensional systems. However, when the system dimension is high, sampling (7) is nontrivial. We aim to address this issue with the proposed framework.

### Safety-critical Control

We consider system (1) with a nominal control policy  $u = \mathcal{N}(x)$ . We define safety of the system as the state staying within a safe set  $\mathcal{C}$ , which is the super-level set of a barrier function  $\phi(x) : \mathbb{R}^n \rightarrow \mathbb{R}$ , i.e.,

$$\mathcal{C} = \{x : \phi(x) \geq 0\}. \quad (8)$$

This definition of safety can characterize a large variety of practical safety requirements (Prajna, Jadbabaie, and Pappas 2007; Ames et al. 2019). Since in stochastic systems safety can only be guaranteed in the sense of probability, we consider long-term safety probability  $F$  of the system defined as below.

**Definition 1** (Safety probability). *Starting from initial state  $x_0 = x \in \mathcal{C}$ , the safety probability  $F$  of system (1) for outlook time horizon  $t$  is defined as the probability of state  $x_\tau$  staying in the safe set  $\mathcal{C}$  over the time interval  $[0, t]$ , i.e.,*

$$F(x, t) = \mathbb{P}(x_\tau \in \mathcal{C}, \forall \tau \in [0, t] \mid x_0 = x). \quad (9)$$




Figure 2: The procedure diagram of the proposed method.

The goal is to find the safety probability  $F$  over the state space for a long-term horizon  $T$ . Once the safety probability is acquired, existing safe control methods can be used to guarantee safety of the system (Wang et al. 2021; Jing and Nakahira 2022). One standard approach to acquire such safety probability is to run Monte Carlo simulation with the nominal controller  $\mathcal{N}$  multiple times and calculate the empirical probability of the system being safe, i.e.,

$$\bar{F}(x, T) = \frac{N_{\text{safe}}}{N} \approx \mathbb{P}(x_\tau \in \mathcal{C}, \forall \tau \in [0, T] \mid x_0 = x), \quad (10)$$

where  $N_{\text{safe}}$  is the number of safe trajectories over  $N$  trajectories. However, such estimation has a sample complexity that scales linearly with horizon  $T$  and exponentially with system dimension  $n$  (Rubino and Tuffin 2009; Wang et al. 2021), and is thus inefficient for high-dimensional systems with long-term safety requirements. We aim to overcome these issues and efficiently estimate long-term safety probabilities of high-dimensional systems with the proposed framework.

## Proposed Method

In the section, we introduce the proposed framework to efficiently estimate value functions and safety probabilities of high-dimensional systems. The method consists of three procedures. We first apply comparison theorem to find low-dimensional features of the system and the associated stochastic processes that characterize their evolution. Then we transform the stochastic process into the solution of certain PDEs. Last, we formulate a physics-informed learning problem to efficiently solve the PDE with special boundary and initial conditions. Additionally, we introduce an autoencoder-like neural network for automatic feature identification. Fig. 2 shows the overall procedures of the proposed framework.

### Comparison Theorem for Feature Identification

We assume the low-dimensional feature can be represented by a smooth function  $p : \mathbb{R}^n \rightarrow \mathbb{R}$ . Here, we use comparison theorem (Ikeda and Watanabe 1977, Theorem 3.1) to find the stochastic process  $\xi$  that describes the exact evolution of  $p(x)$ . We introduce the operator  $\mathcal{A}$  as

$$\mathcal{A}^U(\cdot)(x) = \frac{\partial(\cdot)}{\partial x}(x)f(x) + \frac{\partial(\cdot)}{\partial x}(x)\sigma(x)U + \frac{1}{2} \text{Tr} \left( \frac{\partial^2(\cdot)}{\partial x^2}(x)\sigma(x)\sigma(x)^\top \right), \quad (11)$$

where  $U : \mathbb{R}^n \rightarrow \mathbb{R}^m$  is the control policy that maps the state to the control input, and we define

$$a(x) = \sum_{i,j,k} \sigma_k^i(x) \sigma_k^j(x) \frac{\partial p}{\partial x_i}(x) \frac{\partial p}{\partial x_j}(x), \quad b(x) = \frac{\mathcal{A}^U p(x)}{a(x)}, \quad (12)$$

where  $\sigma_k^i(x)$  is the  $(i, k)$ -th element of the diffusion coefficient  $\sigma(x)$ . For value function estimation, we consider  $\mathcal{A}^0$  for the uncontrolled process where  $U \equiv 0$ , and for safety probability estimation, we consider  $\mathcal{A}^{\mathcal{N}}$  where  $U = \mathcal{N}$  is the nominal control policy. Additionally, let  $\xi \in I \subset \mathbb{R}$  be a scalar variable, where  $I$  is the range of the feature  $p(x)$ . Using similar notation employed in (Ikeda and Watanabe 1977), consider

$$\begin{aligned} a^+(\xi) &= \sup_{x:p(x)=\xi} a(x), & a^-(\xi) &= \inf_{x:p(x)=\xi} a(x) \\ b^+(\xi) &= \sup_{x:p(x)=\xi} b(x), & b^-(\xi) &= \inf_{x:p(x)=\xi} b(x) \end{aligned} \quad (13)$$

**Assumption 2.** We assume that the feature function  $p(x)$  satisfies that  $a^+(\xi) = a^-(\xi) = \alpha(\xi)$  and  $b^+(\xi) = b^-(\xi) = \beta(\xi)$ ,  $\forall \xi \in I$ .

**Assumption 3.** The functions  $\alpha(\xi)$  and  $\beta(\xi)$  are globally Lipschitz continuous in  $\xi \in I \subset \mathbb{R}$ . Moreover,  $a(x) > 0$ ,  $\forall x \in \mathbb{R}^n$ .

**Theorem 4.** Given Assumptions 2 and 3 hold,  $p(x_t)$  with  $x_t$  being sampled from system (1) is characterized by the following stochastic process

$$d\xi_t = \alpha(\xi_t) \beta(\xi_t) dt + \sqrt{\alpha(\xi_t)} d\tilde{B}_t, \quad (14)$$

with  $\xi_0 = p(x_0)$ , and  $\tilde{B}_t$  being a one-dimensional standard Wiener process.

*Proof.* See extended version of the paper (Wang et al. 2023).

Assumption 2 gives the conditions for the upper and lower bounds to match with the actual value  $\alpha$  and  $\beta$ , thus a single stochastic process  $\xi$  can be derived. Until here, we have found the one-dimensional process (14) that characterizes the evolution of the feature of the high-dimensional system without any information loss.

### PDE for Value Function with Feynman-Kac

In this section, we describe how to estimate the value function with the solution of a low-dimensional PDE. We consider  $p(x) = c(x)$ , *i.e.*, the feature of the high-dimensional system is the value of the running cost on state. We set

$$V(x, t) = -\log \varphi(x, t). \quad (15)$$

Let  $\xi = p(x)$ , then from (6) we have

$$\varphi(x, t) = \mathbb{E} \left[ \exp \left( - \int_t^T \xi_\tau d\tau - \xi_T \right) \mid \xi_t = \xi \right] \quad (16)$$

With that, we apply Feynman-Kac formula (Del Moral and Del Moral 2004) on (16) and (14) and get  $\varphi(x, t)$  is the solution of the following two-dimensional PDE

$$\begin{aligned} W_\varphi(\xi, t) &:= \frac{\partial \varphi}{\partial t}(\xi, t) + \alpha(\xi, t) \beta(\xi, t) \frac{\partial \varphi}{\partial \xi}(\xi, t) \\ &+ \frac{1}{2} \alpha(\xi, t) \frac{\partial^2 \varphi}{\partial \xi^2}(\xi, t) - \xi \varphi(\xi, t) = 0, \end{aligned} \quad (17)$$

with initial (terminal) condition

$$\varphi(\xi, T) = \exp\left(-\frac{1}{2}\xi\right). \quad (18)$$

### PDE for Safety Probability

In this section, we describe how to estimate the safety probability with the solution of a low-dimensional PDE. We consider  $p(x) = \phi(x)$ , *i.e.*, the feature of the high-dimensional system is the value of the barrier function for the state. Then we have

$$\begin{aligned} F(x, t) &= \mathbb{P}(x_\tau \in \mathcal{C}, \forall \tau \in [0, t] \mid x_0 = x) \\ &= \mathbb{P}\left(\min_{0 \leq \tau \leq t} \phi(x_\tau) \geq 0\right) = \mathbb{P}\left(\min_{0 \leq \tau \leq t} p(x_\tau) \geq 0\right) \\ &= \mathbb{P}\left(\min_{0 \leq \tau \leq t} \xi_\tau \geq 0\right). \end{aligned} \quad (19)$$

The well-known results on the probability distribution of the first hitting time (Patie and Winter 2008) allow us to obtain (19) as a solution to the two-dimensional PDE given by

$$\begin{aligned} W_F(\xi, t) &:= \frac{\partial F}{\partial t}(\xi, t) - \alpha(\xi, t) \beta(\xi, t) \frac{\partial F}{\partial \xi}(\xi, t) \\ &- \frac{1}{2} \alpha(\xi, t) \frac{\partial^2 F}{\partial \xi^2}(\xi, t) = 0, \end{aligned} \quad (20)$$

with boundary and initial conditions

$$F(\xi, 0) = 1, \quad \xi > 0; \quad F(0, t) = 0, \quad t > 0. \quad (21)$$

### Physic-informed Learning

In this section, we will introduce a physics-informed learning pipeline to solve the PDE for the value function (17) and the PDE for safety probability (20). For conciseness, we will focus on the case of value function estimation, as adaptation to safety probability estimation is trivial where one just needs to replace the variables and the governing PDE.

From (7), we can estimate the value function using path integral control by sampling the uncontrolled process. However, the path integral MC is not sample efficient when we want to know the value function on the entire space, especially for a large horizon  $T$ . Further, efficiently solving the value function PDE (17) using standard numerical methods is challenging.

To leverage the advantages of MC and PDE methods and to overcome their drawbacks, we propose a physics-informed neural network (PINN) to learn the mapping from the feature-time pair to the value function  $\varphi$ . Fig. 3a shows the architecture of the PINN. The PINN takes the feature-time pair  $(\xi, t)$  as the input, and outputs the value function prediction  $\hat{\varphi}$ , the feature and time derivatives  $\frac{\partial \hat{\varphi}}{\partial \xi}$  and  $\frac{\partial \hat{\varphi}}{\partial t}$ , and the Hessian  $\frac{\partial^2 \hat{\varphi}}{\partial \xi^2}$ , which come naturally from the automatic differentiation in deep learning frameworks such as PyTorch (Paszke et al. 2019) and TensorFlow (Abadi et al. 2016). Assume the PINN is parameterized by  $\theta$ , the loss function is defined as

$$\mathcal{L}(\theta) = \omega_p \mathcal{L}_p(\theta) + \omega_d \mathcal{L}_d(\theta), \quad (22)$$

where

$$\begin{aligned}\mathcal{L}_p(\theta) &= \frac{1}{|\mathcal{P}|} \sum_{(\xi, t) \in \mathcal{P}} \|W_{\hat{\varphi}_\theta}(\xi, t)\|_2^2, \\ \mathcal{L}_d(\theta) &= \frac{1}{|\mathcal{D}|} \sum_{(\xi, t) \in \mathcal{D}} \|\hat{\varphi}_\theta(\xi, t) - \bar{\varphi}(\xi, t)\|_2^2.\end{aligned}\quad (23)$$

Here,  $\bar{\varphi}$  is the training data,  $\hat{\varphi}_\theta$  is the prediction from the PINN,  $\mathcal{P}$  and  $\mathcal{D}$  are the training point sets for the physics model and external data, respectively. The loss function  $\mathcal{L}$  consists of two parts, the physics model loss  $\mathcal{L}_p$  and data loss  $\mathcal{L}_d$ . The physics model loss  $\mathcal{L}_p$  measures the satisfaction of the PDE constraints for the learned output. It calculates the actual PDE equation value  $W_{\hat{\varphi}_\theta}$ , which is supposed to be 0, and use its 2-norm as the loss. The data loss  $\mathcal{L}_d$  measures the accuracy of the prediction of PINN on the training data. It calculates the mean square error between the PINN prediction and the training data point as the loss. The overall loss function  $\mathcal{L}$  is the weighted sum of the physics model loss and data loss with weighting coefficients  $\omega_p$  and  $\omega_d$ . Though out of the scope of this paper, theoretical analysis on the approximation error of PINNs and neural operators can be found in (Wang and Nakahira 2023; Lu et al. 2021a; Kovachki, Lanthaler, and Mishra 2021).

**Remark 5.** For finding the optimal control, one can further enhance the sample complexity using path integral control together with value function prediction from the PINN. Denote  $\hat{V}$  as the PINN value function prediction, we can initially estimate the optimal control from (5) with  $\hat{u}(x, t) = -\sigma(x)^\top \nabla \hat{V}(x, t)$ . Then we can refine the optimal control using importance sampling (Thijssen and Kappen 2015) with the following procedure

$$\begin{aligned}u^*(x, t) - \hat{u}(x, t) &= \\ \lim_{s \searrow t} \frac{\mathbb{E}^P \left\{ \exp \left\{ -S^{\hat{u}}(t) \right\} \int_t^s dW_\tau \mid x_t = x \right\}}{(s-t) \mathbb{E}^P \left\{ \exp \left\{ -S^{\hat{u}}(t) \right\} \mid x_t = x \right\}}\end{aligned}\quad (24)$$

where  $P$  is the process with regard to  $\hat{u}$  and

$$S^{\hat{u}}(t) = \int_t^T w(x_\tau, \hat{u}_\tau) d\tau + \int_t^T \hat{u}_\tau^\top dW_\tau + c(x_T). \quad (25)$$

Essentially, one can use the sampled cost function (25) with control policy  $\hat{u}$  to estimate the optimal control policy  $u^*$  with (24). The sample complexity for estimating the expectation in (24) is much lower than the naive estimation with (7) and (5) due to the fact that path generated from  $\hat{u}$  is much closer to the optimal path. The importance sampling theory provides theoretical analysis on the improvement of the sample complexity (Thijssen and Kappen 2015).

### Generalization: Arbitrary Feature Dimension

In this section we generalize the previous results such that the representation of the system can be of arbitrary dimension. The procedure consists of two key steps. First, we use comparison theorem to find a multidimensional representation of the value function and the associated multidimensional process, i.e.,  $\xi = [\xi^{(1)}, \xi^{(2)}, \dots, \xi^{(k)}]^\top$  where  $k$  is

the dimension of the reduced representation. Then, we apply the high-dimensional Feynman-Kac formula (Pham 2009, Theorem 1.3.17) to transform the stochastic process  $\xi$  to a  $k$ -dimensional PDE which can be solved by the PINN.

For the first step, we find functions  $p(x) = [p_1(x), p_2(x), \dots, p_k(x)]^\top$  as the low dimensional representation of the original system. We define  $a_i^{+-}$  and  $b_i^{+-}$  similar to (13) and assume Assumption 2 and 3 hold for  $\forall i$ . Then from the comparison theorem, we can find stochastic processes

$$d\xi_t^{(i)} = \alpha_i(\xi_t^{(i)}) \beta_i(\xi_t^{(i)}) dt + \sqrt{\alpha_i(\xi_t^{(i)})} d\tilde{B}_t^{(i)} \quad (26)$$

for  $i = 1, 2, \dots, k$  that characterize  $p_i(x)$ , where  $\tilde{B}_t^{(i)}$  is one-dimensional standard Wiener processes.

**Assumption 6.** For value function estimation, we assume the running-cost can be represented by the following function

$$c(x) = r(\xi) = r(p_1(x), p_2(x), \dots, p_k(x)) \quad (27)$$

where  $r: \mathbb{R}^k \rightarrow \mathbb{R}$  is a continuous function.

Assume Assumption 6 holds, then

$$\varphi(x, t) = \mathbb{E} \left[ \exp \left( - \int_t^T r(\xi_\tau) d\tau - r(\xi_T) \right) \mid \xi_t = \xi \right]. \quad (28)$$

From the high-dimensional Feynman-Kac formula (Pham 2009, Theorem 1.3.17), we have  $\varphi$  as the solution of the following PDE

$$\begin{aligned}r\varphi - \frac{\partial \varphi}{\partial t} - \mathcal{G}_t \varphi &= 0, \quad \text{on } \mathbb{R}^k \times [0, T) \\ \varphi(\cdot, T) &= \exp(-r(\cdot)), \quad \text{on } \mathbb{R}^k,\end{aligned}\quad (29)$$

where  $r$  is given by (27) and

$$\mathcal{G}_t(\cdot) = \alpha(\xi_t) \beta(\xi_t) \cdot \frac{\partial(\cdot)}{\partial \xi} + \frac{1}{2} \text{Tr} \left( \mathbf{a}(\xi_t) \frac{\partial^2(\cdot)}{\partial \xi^2} \right), \quad (30)$$

with

$$\alpha(\xi_t) \beta(\xi_t) = \begin{bmatrix} \alpha_1(\xi_t^{(1)}) \beta_1(\xi_t^{(1)}) \\ \alpha_2(\xi_t^{(2)}) \beta_2(\xi_t^{(2)}) \\ \vdots \\ \alpha_k(\xi_t^{(k)}) \beta_k(\xi_t^{(k)}) \end{bmatrix}, \quad (31)$$

and

$$\mathbf{a}(\xi_t) = \begin{bmatrix} \alpha_1(\xi_t) & & & \\ & \alpha_2(\xi_t) & & \\ & & \ddots & \\ & & & \alpha_k(\xi_t) \end{bmatrix}. \quad (32)$$

**Assumption 7.** Similarly, for safety probability estimation, we assume the barrier function can be represented by the following function

$$\phi(x) = r(\xi) = r(p_1(x), p_2(x), \dots, p_k(x)), \quad (33)$$

where  $r: \mathbb{R}^k \rightarrow \mathbb{R}$  is a continuous function.

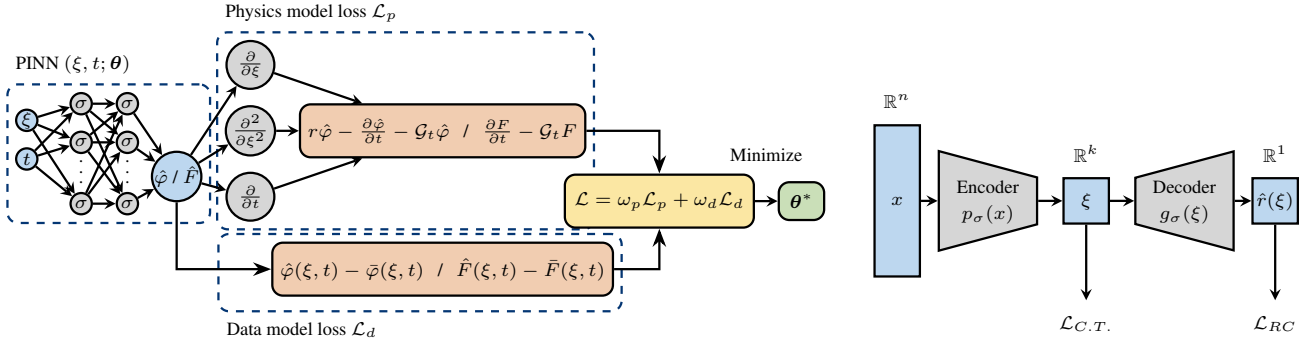


Figure 3: (a) The training scheme of the physics-informed neural network (PINN). (b) Autoencoder-like network architecture.

Then the safety probability can be written as

$$F(x, t) = \mathbb{P} \left( \min_{0 \leq \tau \leq t} \phi(x_\tau) \geq 0 \right) = \mathbb{P} \left( \min_{0 \leq \tau \leq t} r(\xi_\tau) \geq 0 \right). \quad (34)$$

We define  $\mathcal{B} = \{\xi : r(\xi) \geq 0\}$ . From the probability distribution of hitting time (Patie and Winter 2008) we have that (34) can be characterized by the solution of the following PDE

$$\frac{\partial F}{\partial t} - \mathcal{G}_t F = 0, \text{ on } [0, T) \times \mathbb{R}^n \quad (35)$$

$$F(\xi, 0) = 1, \xi \in \mathcal{B}; \quad F(\xi, t) = 0, \xi \in \partial \mathcal{B}. \quad (36)$$

**Remark 8.** *The Assumptions 6 or 7 where the running-cost can be represented by (27) or the barrier function can be represented by (33) is a necessary condition for the proposed method to work. The high-dimensional system must admit a low-dimensional representation of its value function/safety probability.*

### Deep Learning for Feature Identification

For high dimensional systems with complex structure, deriving feature maps  $p_1, p_2, \dots, p_k$  such that Assumptions 2 and 3 hold is a challenging problem. Thus, we propose an autoencoder-like neural network (Fig. 3b) to automatically identify lower-dimensional features that meet the bounding requirements of Assumption 2 and sufficiently represent the cost/barrier function (Remark 8). The network takes an input state  $x \in \mathbb{R}^n$  and outputs a low-dimensional representation  $\xi \in \mathbb{R}^k$  via the encoder  $p_\sigma(x)$  and the function  $\hat{r}(\xi) \in \mathbb{R}$  via the decoder  $g_\sigma(\xi)$ . We use  $\theta$  as the parameters of the autoencoder-like model. The loss function is defined as

$$\mathcal{L}_{AE}(\theta) = w_{RC} \mathcal{L}_{RC}(\theta) + w_{C.T.} \mathcal{L}_{C.T.}(\theta) \quad (37)$$

where

$$\mathcal{L}_{RC}(\theta) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} (c(x) - \hat{r}(\xi; \theta))^2, \quad (38)$$

$$\mathcal{L}_{C.T.}(\theta) = \frac{1}{k} \sum_{i=1}^k \frac{1}{|\mathcal{R}_i|} \sum_{\xi \in \mathcal{R}_i} \frac{1}{|\mathcal{M}_{\xi, i}|} \sum_{x \in \mathcal{M}_{\xi, i}} \dots \|\nabla_x a_i(x; \theta)\|_2^2 + \|\nabla_x b_i(x; \theta)\|_2^2 \quad (39)$$

Here,  $\mathcal{X}$  is the discretized state-space,  $k$  is the dimension of the feature space,  $\mathcal{R}_i$  is the range of  $p_i$  and  $\mathcal{M}_{\xi, i} = \{x : p_i(x; \theta) = \xi\}$  is the set of  $x$  that maps to the same value of  $\xi \in \mathcal{R}_i$ . The reconstruction loss  $\mathcal{L}_{RC}$  measures how well the reconstruction  $\hat{r}(\xi)$  represents  $c(x)$  or  $\phi(x)$ , corresponding to Remark 8. The comparison theorem loss  $\mathcal{L}_{C.T.}$  enforces the condition that  $a_i(x)$  and  $b_i(x)$  are constant for all  $x \in \mathcal{M}_{\xi, i}$ , for each  $\xi \in \mathcal{R}_i$ , and for each feature  $p_1, \dots, p_k$ . This is a sufficient condition to achieve Assumption 2. The overall loss function is the weighted sum of the reconstruction loss and comparison theorem loss, where weights are chosen according to the desired strength and balance on the cost/barrier function reconstruction and the satisfaction of the comparison theorem. We refer readers to extended version of the paper (Wang et al. 2023) for algorithm details.

### Experiments

In this section, we show experiment results of the proposed method with both qualitative and quantitative analysis.

#### Value Function Estimation

We consider a 1000-dimensional system for which we want a 2-dimensional representation of the optimal value function. The system dynamics is given by

$$dx = \bar{A}xdt + \sigma(udt + dw), \quad (40)$$

where  $x \in \mathbb{R}^{1000}$  is the state,  $u \in \mathbb{R}^{1000}$  in the control, and  $dw$  is the 1000-dimensional standard Wiener process. We set  $\sigma = I_{1000}$  to be the identity matrix, and set  $\bar{A} = \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix}$  where  $A \in \mathbb{R}^{500 \times 500}$ . Let  $a_{i,j}$  be the entry of  $A$  at  $i$ -th row and  $j$ -th column. We choose  $A$  such that  $a_{i,i} = 1.1$ ,  $a_{i,(i+2)|500} = a_{i,(i+4)|500} = 0.1$  and  $a_{i,(i+6)|500} = a_{i,(i+8)|500} = -0.1$  for  $\forall i = 1, 2, \dots, 500$ , where  $|$  is the mod operator. The running cost function is assumed to be

$$c(x) = \frac{1}{500} \left( \sum_{i=1}^{500} x_i \right)^2 + \frac{1}{500} \left( \sum_{j=501}^{1000} x_j \right)^2. \quad (41)$$

We pick two features of the state,

$$\xi_1 = p_1(x) = \sum_{i=1}^{500} x_i, \quad \xi_2 = p_2(x) = \sum_{j=501}^{1000} x_j. \quad (42)$$



Then from (13), we have

$$\alpha_1(\xi_1) = \alpha_2(\xi_2) = 500; \beta_1(\xi_1) = \frac{\xi_1}{500}, \beta_2(\xi_2) = \frac{\xi_2}{500}, \quad (43)$$

thus satisfying Assumptions 2 and 3, and the running cost function can be written as

$$r(\xi) = \frac{1}{500}\xi_1^2 + \frac{1}{500}\xi_2^2, \quad (44)$$

thus satisfying Assumption 6. By the Feynman-Kac formula, we know that the exponential of the optimal value function  $\varphi$  is the solution of the following PDE

$$0 = r(\xi)\mu - \frac{\partial\mu}{\partial t} - \alpha(\xi)\beta(\xi)\frac{\partial\mu}{\partial\xi} - \frac{1}{2}\text{Tr}\left(\mathbf{a}(\xi)\frac{\partial^2\mu}{\partial\xi^2}\right) \quad (45)$$

$$= \frac{\xi_1^2 + \xi_2^2}{500}\mu - \frac{\partial\mu}{\partial t} - \xi_1\frac{\partial\mu}{\partial\xi_1} - \xi_2\frac{\partial\mu}{\partial\xi_2} - \frac{\partial^2\mu}{\partial\xi_1^2} - 250\frac{\partial^2\mu}{\partial\xi_2^2}$$

$$\mu(\xi, T) = \exp\left(-\frac{1}{500}\xi_1^2 - \frac{1}{500}\xi_2^2\right). \quad (46)$$

With that, we generate data for  $\bar{\varphi}(\xi, t)$  on spatial-temporal space  $\Omega \times \mathbb{T} = [1, 2]^2 \times [0, 1.5]$ , with grid size  $d\xi = 0.1$  and  $dt = 0.1$  and train the PINN on  $\Omega \times \mathbb{T} = [1, 2]^2 \times [1, 1.5]$ . We use a PINN with 3 hidden layers and 32 neurons per layer to learn the value function  $\varphi$ . The activation function is chosen as hyperbolic tangent function (tanh). We use the Adam optimizer (Kingma and Ba 2014) for training with initial learning rate set as 0.001. The PINN parameters  $\theta$  are initialized via Glorot uniform initialization and the weights in the loss function (22) are set to be  $\omega_p = \omega_d = 1$ . The simulation is constructed based on the DeepXDE framework (Lu et al. 2021b). Fig. 4 shows the estimated value function from the path integral MC and the proposed method. It can be seen that the proposed framework is able to estimate value functions accurately, while the path integral method has significantly more noise. Also, the computation time for training the PINN is significantly less than sampling path integral MC (80s v.s  $\sim 3000$ s). Note that the PINN is able to estimate the value function at unseen regions in the state space and generalize to longer time horizons, as the testing data at  $t = 0.5$  is not seen by the PINN during training. We refer readers to the extended version of the paper (Wang et al. 2023) for the safety probability estimation setting where similar results can be obtained.

### Sample Complexity

We further examine the computation complexity of the proposed method to show its advantages in sample efficiency. We consider the problem of value function estimation of a 3-dimensional system with similar dynamics and cost defined in (40) and (41). See extended version of the paper (Wang et al. 2023) for details of the setting. Fig. 5 shows the percentage error of the estimated value function with different number of samples for path integral MC and PINN with and without using comparison theorem for dimension reduction. The path integral MC with dimension reduction has sample complexities that are one order less than MC without dimension reduction, which indicates the efficacy of the dimension reduction scheme from the comparison theorem. The




Figure 4: Estimation of the exponential of value function at  $t = 0.5$  for the 1000-dimensional system by path integral MC (left), and by the proposed method (right).




Figure 5: Percentage error of the estimated value function with path integral MC.

PINN further reduces sample complexity and achieves the best trade-off between accuracy and computation.

### Feature Learning

We verify our autoencoder-like network by learning the features and cost function (similar to equations (42) and (44)) for the 3-dimensional system used for sample complexity analysis. The network consists of 5 fully connected hidden layers of sizes 100, 10, 2, 10, 100, respectively, with the activation function as hyperbolic tangent. We use the Adam optimizer with initial learning rate set as 0.001. The parameters of the network are initialized via the Glorot uniform initialization. The weights in the loss function (37) are set as  $w_{RC} = 1, w_{C.T.} = 10$ . The network is trained on a  $[0, 1]^3$  state-space with grid size 0.01. The network successfully identified the features derived analytically with  $\text{MSE}(\xi_1, \hat{\xi}_1) = 0.2$  and  $\text{MSE}(\xi_2, \hat{\xi}_2) = 0.06$ .

### Conclusion

We propose a unified framework for value function and safety probability estimation of high-dimensional stochastic systems. The novel dimensionality reduction technique uses the comparison theorem to generate low-dimensional stochastic processes that provide an exact characterization of the cost/barrier function, significantly improving sample complexity. We then transform the low-dimensional process into a low-dimensional PDE, and leverage physics-informed learning to generalize solutions into longer time horizons and unseen regions of state-space. We also achieve automatic feature identification through a specially designed autoencoder-like neural network. Experiment results show the efficacy of the proposed method. Future work includes application to multi-agent robotic control systems.

## Acknowledgments

This project is funded in part by Carnegie Mellon University's Mobility21 National University Transportation Center, which is sponsored by the US Department of Transportation, in part by JST, PRESTO Grant Number JPMJPR2136, Japan, in part by the Department of the Navy, Office of Naval Research, grant number N00014-23-1-2252, in part by ACT-X, Japan, under Grant JPMJAX210L, and in part by AFOSR Grant FA9550-20-1-0101. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research.

## References

- Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G. S.; Davis, A.; Dean, J.; Devin, M.; et al. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- Ames, A. D.; Coogan, S.; Egerstedt, M.; Notomista, G.; Sreenath, K.; and Tabuada, P. 2019. Control barrier functions: Theory and applications. In *2019 18th European control conference (ECC)*, 3420–3431. IEEE.
- Bansal, S.; Bajcsy, A.; Ratner, E.; Dragan, A. D.; and Tomlin, C. J. 2020. A hamilton-jacobi reachability-based framework for predicting and analyzing human motion for safe planning. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 7149–7155. IEEE.
- Bertini, L.; Faggionato, A.; and Gabrielli, D. 2015. Large deviations of the empirical flow for continuous time Markov chains. *arXiv:1210.2004*.
- Botev, Z. I.; L'Ecuyer, P.; and Tuffin, B. 2013. Markov chain importance sampling with applications to rare event probability estimation. *Statistics and Computing*, 23(2): 271–285.
- Boué, M.; and Dupuis, P. 1998. A variational representation for certain functionals of Brownian motion. *The Annals of Probability*, 26(4): 1641–1659.
- Bressloff, P. C.; and Newby, J. M. 2014. Path integrals and large deviations in stochastic hybrid systems. *Physical Review E*, 89(4): 042701.
- Cai, S.; Mao, Z.; Wang, Z.; Yin, M.; and Karniadakis, G. E. 2022. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mechanica Sinica*, 1–12.
- Carius, J.; Ranftl, R.; Farshidian, F.; and Hutter, M. 2022. Constrained stochastic optimal control with learned importance sampling: A path integral approach. *The International Journal of Robotics Research*, 41(2): 189–209.
- Cheng, R.; Khojasteh, M. J.; Ames, A. D.; and Burdick, J. W. 2020. Safe multi-agent interaction through robust control barrier functions with learned uncertainties. In *2020 59th IEEE Conference on Decision and Control (CDC)*, 777–783. IEEE.
- Chern, A.; Wang, X.; Iyer, A.; and Nakahira, Y. 2021. Safe control in the presence of stochastic uncertainties. *arXiv preprint arXiv:2104.01259*.
- Clark, A. 2019. Control barrier functions for complete and incomplete information stochastic systems. In *2019 American Control Conference (ACC)*, 2928–2935. IEEE.
- Del Moral, P.; and Del Moral, P. 2004. *Feynman-kac formulae*. Springer.
- Feng, R. 2014. A comparative study of risk measures for guaranteed minimum maturity benefits by a PDE method. *North American Actuarial Journal*, 18(4): 445–461.
- Fleming, W. H.; and Soner, H. M. 2006. *Controlled Markov processes and viscosity solutions*, volume 25. Springer Science & Business Media.
- Frankowska, H.; and Zhang, X. 2020. Necessary conditions for stochastic optimal control problems in infinite dimensions. *Stochastic Processes and their Applications*, 130(7): 4081–4103.
- Gómez, V.; Thijssen, S.; Symington, A.; Hailes, S.; and Kappen, H. 2016. Real-time stochastic optimal control for multi-agent quadrotor systems. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 26, 468–476.
- Gorodetsky, A.; Karaman, S.; and Marzouk, Y. 2018. High-dimensional stochastic optimal control using continuous tensor decompositions. *The International Journal of Robotics Research*, 37(2-3): 340–377.
- Han, J.; Jentzen, A.; and Weinan, E. 2018. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34): 8505–8510.
- Hanna, J. P.; Niekum, S.; and Stone, P. 2021. Importance sampling in reinforcement learning with an estimated behavior policy. *Machine Learning*, 110(6): 1267–1317.
- Hewing, L.; and Zeilinger, M. N. 2018. Stochastic model predictive control for linear systems using probabilistic reachable sets. In *2018 IEEE Conference on Decision and Control (CDC)*, 5182–5188. IEEE.
- Huang, X.; Chen, J.; and Zhu, H. 2016. Assessing small failure probabilities by AK-SS: An active learning method combining Kriging and Subset Simulation. *Structural Safety*, 59: 86–95.
- Huh, S.; and Yang, I. 2020. Safe reinforcement learning for probabilistic reachability and safety specifications: A Lyapunov-based approach. *arXiv preprint arXiv:2002.10126*.
- Ikeda, N.; and Watanabe, S. 1977. A comparison theorem for solutions of stochastic differential equations and its applications. *Osaka Journal of Mathematics*, 14: 619–633.
- Jing, H.; and Nakahira, Y. 2022. Probabilistic safety certificate for multi-agent systems. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, 5343–5350. IEEE.
- Kappen, H. J. 2005. Path integrals and symmetry breaking for optimal control theory. *Journal of statistical mechanics: theory and experiment*, 2005(11): P11011.
- Katsoulakis, M. A.; and Vilanova, P. 2020. Data-driven, variational model reduction of high-dimensional reaction networks. *Journal of Computational Physics*, 401: 108997.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kovachki, N.; Lanthaler, S.; and Mishra, S. 2021. On universal approximation and error bounds for Fourier neural operators. *The Journal of Machine Learning Research*, 22(1): 13237–13312.
- Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; and Karniadakis, G. E. 2021a. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3): 218–229.
- Lu, L.; Meng, X.; Mao, Z.; and Karniadakis, G. E. 2021b. DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1): 208–228.
- Madhusani, U.; Dey, B.; Leonard, N.; and Chakraborty, A. 2021. Hamiltonian Q-Learning: Leveraging Importance-sampling for Data Efficient {RL}.



- Meng, Y.; and Liu, J. 2022. Sufficient conditions for robust probabilistic reach-avoid-stay specifications using stochastic Lyapunov-barrier functions. In *2022 American Control Conference (ACC)*, 2283–2288. IEEE.
- Misyris, G. S.; Venzke, A.; and Chatzivasileiadis, S. 2020. Physics-informed neural networks for power systems. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, 1–5. IEEE.
- Nishimura, Y.; and Hoshino, K. 2023. Control Barrier Functions for Stochastic Systems With Quantitative Evaluation of Probability. [arXiv:2209.08728](https://arxiv.org/abs/2209.08728).
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Patie, P.; and Winter, C. 2008. First exit time probability for multidimensional diffusions: A PDE-based approach. *Journal of computational and applied mathematics*, 222(1): 42–53.
- Patil, A.; Duarte, A.; Smith, A.; Bisetti, F.; and Tanaka, T. 2022. Chance-constrained stochastic optimal control via path integral and finite difference methods. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, 3598–3604. IEEE.
- Pereira, M.; Wang, Z.; Exarchos, I.; and Theodorou, E. 2021. Safe optimal control using stochastic barrier functions and deep forward-backward sdes. In *Conference on Robot Learning*, 1783–1801. PMLR.
- Pham, H. 2009. *Continuous-time stochastic control and optimization with financial applications*, volume 61. Springer Science & Business Media.
- Prajna, S.; Jadbabaie, A.; and Pappas, G. J. 2007. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Transactions on Automatic Control*, 52(8): 1415–1428.
- Raissi, M.; Perdikaris, P.; and Karniadakis, G. E. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378: 686–707.
- Rashki, M. 2021. SESC: A new subset simulation method for rare-events estimation. *Mechanical Systems and Signal Processing*, 150: 107139.
- Rubino, G.; and Tuffin, B. 2009. *Rare event simulation using Monte Carlo methods*. John Wiley & Sons.
- Sahli Costabal, F.; Yang, Y.; Perdikaris, P.; Hurtado, D. E.; and Kuhl, E. 2020. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8: 42.
- Samuelson, S.; and Yang, I. 2018. Safety-aware optimal control of stochastic systems using conditional value-at-risk. In *2018 Annual American Control Conference (ACC)*, 6285–6290. IEEE.
- Santoyo, C.; Dutreix, M.; and Coogan, S. 2021. A barrier function approach to finite-time stochastic system verification and control. *Automatica*, 125: 109439.
- Satoh, S.; and Kappen, H. J. 2020. Nonlinear stochastic optimal control with input saturation constraints based on path integrals. *IEEJ Transactions on Electrical and Electronic Engineering*, 15(8): 1169–1175.
- Satoh, S.; Kappen, H. J.; and Saeki, M. 2016. An iterative method for nonlinear stochastic optimal control based on path integrals. *IEEE Transactions on Automatic Control*, 62(1): 262–276.
- Stadie, B. C.; Yang, G.; Houthoofd, R.; Chen, X.; Duan, Y.; Wu, Y.; Abbeel, P.; and Sutskever, I. 2018. Some considerations on learning to explore via meta-reinforcement learning. [arXiv preprint arXiv:1803.01118](https://arxiv.org/abs/1803.01118).
- Theodorou, E.; Buchli, J.; and Schaal, S. 2010. A generalized path integral control approach to reinforcement learning. *The Journal of Machine Learning Research*, 11: 3137–3181.
- Theodorou, E. A.; and Todorov, E. 2012. Relative entropy and free energy dualities: Connections to path integral and kl control. In *2012 IEEE 51st IEEE conference on decision and control (cdc)*, 1466–1473. IEEE.
- Thijssen, S.; and Kappen, H. 2015. Path integral control and state-dependent feedback. *Physical Review E*, 91(3): 032104.
- Wang, Z.; Jing, H.; Kurniawan, C.; Chern, A.; and Nakahira, Y. 2021. Myopically Verifiable Probabilistic Certificate for Long-term Safety. [arXiv preprint arXiv:2110.13380](https://arxiv.org/abs/2110.13380).
- Wang, Z.; Keller, R.; Deng, X.; Hoshino, K.; Tanaka, T.; and Nakahira, Y. 2023. Physics-informed representation and learning: Control and risk quantification. [arXiv preprint arXiv:2312.10594](https://arxiv.org/abs/2312.10594).
- Wang, Z.; and Nakahira, Y. 2023. A Generalizable Physics-informed Learning Framework for Risk Probability Estimation. In *Learning for Dynamics and Control Conference*, 358–370. PMLR.
- Williams, G.; Aldrich, A.; and Theodorou, E. A. 2017. Model predictive path integral control: From theory to parallel computation. *Journal of Guidance, Control, and Dynamics*, 40(2): 344–357.
- Williams, G.; Wagener, N.; Goldfain, B.; Drews, P.; Rehg, J. M.; Boots, B.; and Theodorou, E. A. 2017. Information theoretic MPC for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 1714–1721. IEEE.
- Yaghoubi, S.; Majd, K.; Fainekos, G.; Yamaguchi, T.; Prokhorov, D.; and Hoxha, B. 2020. Risk-bounded control using stochastic barrier functions. *IEEE Control Systems Letters*, 5(5): 1831–1836.
- Yamada, T. 1973. On a comparison theorem for solutions of stochastic differential equations and its applications. *Journal of Mathematics of Kyoto University*, 13(3): 497 – 512.
- Zhao, Y.; and Wang, Z. 2022. Subset simulation with adaptable intermediate failure probability for robust reliability analysis: An unsupervised learning-based approach. *Structural and Multidisciplinary Optimization*, 65(6): 172.