

A Generalizable Physics-informed Learning Framework for Risk Probability Estimation

Zhuoyuan Wang

ZHUOYUAW@ANDREW.CMU.EDU

Yorie Nakahira

YORIE@CMU.EDU

Electrical and Computer Engineering Department, Carnegie Mellon University, USA

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

Accurate estimates of long-term risk probabilities and their gradients are critical for many stochastic safe control methods. However, computing such risk probabilities in real-time and in unseen or changing environments is challenging. Monte Carlo (MC) methods cannot accurately evaluate the probabilities and their gradients as an infinitesimal divisor can amplify the sampling noise. In this paper, we develop an efficient method to evaluate the probabilities of long-term risk and their gradients. The proposed method exploits the fact that long-term risk probability satisfies certain partial differential equations (PDEs), which characterize the neighboring relations between the probabilities, to integrate MC methods and physics-informed neural networks. Numerical results show the proposed method has better sample efficiency, generalizes well to unseen regions, and can adapt to systems with changing parameters. The proposed method can also accurately estimate the gradients of risk probabilities, which enables first- and second-order techniques on risk probabilities to be used for learning and control.

Keywords: Stochastic safe control; physics-informed learning; risk probability estimation.

1. Introduction

Safe control for stochastic systems is important yet a key challenge for deploying autonomous systems in the real world. In the past decades, many stochastic control methods have been proposed to ensure safety of systems with noises and uncertainties, including stochastic reachabilities (Prandini and Hu, 2008), Condition Value-at-Risk (Ahmadi et al., 2021), chance-constrained predictive control (Nakka et al., 2020) and *etc.* Despite the huge amount of stochastic safe control methods, many of them rely on accurate estimates of long-term risk probabilities or their gradients to guarantee long-term safety (Abate et al., 2008; Chapman et al., 2019; Santoyo et al., 2021; Wang et al., 2021). To get such accurate estimates is non-trivial, and here we list the challenges.

- *High computation complexity.* Estimating long-term risk is computationally expensive, because the possible state trajectories scale exponentially with regard to the time horizon. Besides, the values of risk probability are often small in safety-critical systems, thus huge amounts of sample trajectories are needed to capture the rare unsafe event (Janssen, 2013).
- *Sample inefficiency.* Generalization of risk estimation to the full state space is hard to achieve for sample-based methods, since each point of interest requires one separate simulation. The sample complexity increases linearly with respect to the number of points for evaluation. Besides, most of the existing methods require re-evaluation of the risk probability for any changes of system parameters, which further degrades sample efficiency (Zuev, 2015).

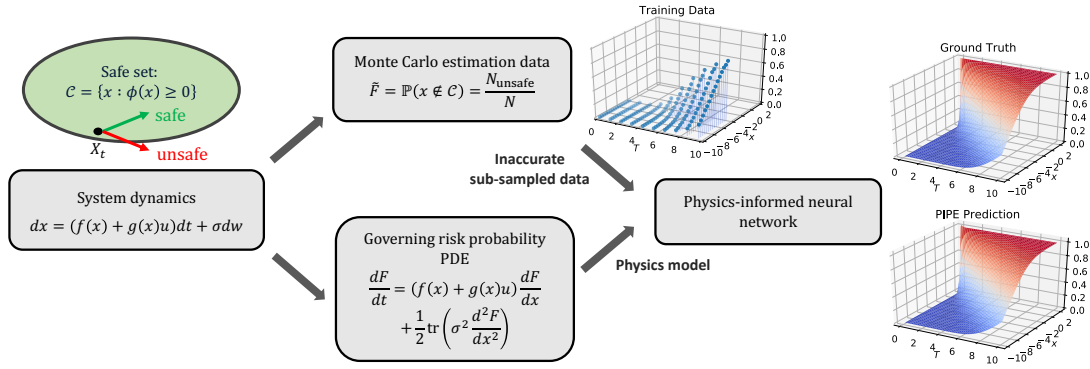


Figure 1: The overview diagram of the proposed PIPE framework. The system of interest is in the form of (1) with safe set defined as (2). For the training data, one can acquire the empirical risk probabilities by forward simulating the system dynamics and calculating the ratio of unsafe trajectories over all trajectories. For the physics model, we know that the mapping between state-time pair to the risk probability satisfies a governing convection diffusion equation (Theorem 3). The PIPE framework uses physics-informed neural networks to learn the risk probability by taking the empirical data for training and use the physics model as constraints. The resulting PIPE will give more accurate prediction on risk probabilities in a more sample efficient way than Monte Carlo or its variants, and can generalize to unseen regions in the state space and unknown parameters in the system dynamics thanks to its integration of data and physics models.

- *Lack of direct solutions.* For control affine systems, recent study (Chern et al., 2021) suggests that the risk probability can be characterized by the solution of partial differential equations (PDEs). It provides an analytical way to calculate the risk probability, but to solve the time-varying PDE and get the actual probability value is non-trivial.
- *Noisy estimation of probability gradients.* Estimating gradients of risk probabilities is difficult, as sampling noise is amplified by the infinitesimal divisor while taking derivatives over the estimated risk probabilities.

To resolve the abovementioned issues, based on study (Chern et al., 2021) and inspired by recently advanced physics-informed neural networks, we propose a physics-informed learning framework PIPE (Physics-Informed Probability Estimator) to estimate risk probabilities in an efficient and generalizable manner. Fig. 1 shows the overview diagram of the proposed PIPE framework. The framework takes both data and physics models into consideration, and by combining the two, we achieve better sample efficiency and the ability to generalize to unseen regions in the state space and unknown parameters in the system dynamics. The use of deep neural networks enables the efficient learning of complex PDEs, and the consideration of physics model further enhances the efficiency as it allows imperfect noisy data for training. The resulting framework takes only inaccurate sample data in a sparse sub-region of the state space for training, and is able to accurately predict the risk probability over the whole state space for systems with different parameters.

2. Related Works

2.1. Stochastic safe control

Stochastic safe control is a heated topic in recent years, as safety under uncertainties and noises becomes the key challenge of many real-world autonomous systems. Stochastic reachability analysis takes the stochastic dynamics model and forward rollouts the possible trajectories to get the safety probability of any given state, and use this information to design suitable safe controllers (Abate et al., 2008; Prandini and Hu, 2008). Conditional Value-at-Risk considers the risk measure of the system and guarantees the expected risk value to always decrease conditioned on the previous states of the system, and thus guarantees safety (Samuelson and Yang, 2018; Singletary et al., 2022). Chance-constrained predictive control takes probabilistic safety requirements as the constraints in an optimization-based controller, and solves a minimal distance control to a nominal controller to find its safe counterpart (Nakka et al., 2020; Zhu and Alonso-Mora, 2019; Pfrommer et al., 2022). While these methods provide theoretical guarantees on safety, all of them require accurate estimation of risk probabilities or their gradients to yield desirable performance, and to get accurate estimates itself challenging. We tackle this problem by combining physics models and data to provide accurate estimates of risk probability and its gradient.

2.2. Rare event simulation

Rare event simulation considers the problem of estimating the probability of rare events in the system, and is highly related to risk probability estimation because the risk probability is often small in safety-critical systems. Here, we list a few widely adopted approaches for rare event simulation. Standard MC forward runs the system dynamics multiple times to empirically estimate the risk probability by calculating the unsafe trajectory numbers over the total trajectory number (Rubino and Tuffin, 2009). Standard MC is easy to implement, but needs huge amounts of sample trajectories to get accurate estimation, which becomes impractical when the required accuracy is high. Importance sampling methods calculate the risk probability on a shifted new distribution to improve the sample efficiency, but needs good prior information on the re-sampled distribution for reasonable performance enhancement, which is hard to achieve for complex systems (C erou et al., 2012; Botev et al., 2013). Subset simulation calculates the risk probability conditioned on intermediate failure events that are easier to estimate, to further improve sample efficiency (Au and Beck, 2001). However, computation efficiency remains an issue and generalization to the entire state space is hard to achieve, as the estimation can only be conducted at a single point once (Zuev, 2015). There are no known methods that can compute the risk probability in an integrated way for the entire state space, and to do so with high sample efficiency. We address this problem by proposing a learning framework that considers both data and model to give generalizable prediction results with high sample efficiency.

2.3. 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 partial differential equations (Raissi et al., 2019). PINN takes both data and the physics model of the system into account, and is able to solve the forward problem of getting PDE solutions, and the inverse problem of discovering underlying governing PDEs from data. PINN has been widely used in power

systems (Misyris et al., 2020), fluid mechanics (Cai et al., 2022) and medical care (Sahli Costabal et al., 2020). For stochastic safe control, previous works use PINN to solve the initial value problem of deep backward stochastic differential equations to derive an end-to-end myopic safe controller (Han et al., 2018; Pereira et al., 2021). However, to our best knowledge there is no work that considers PINN for risk probability estimation, especially on the full state-time space scale. In this work we take the first step towards leveraging PINN on the problem of risk probability estimation.

3. Problem Formulation

We consider a control system with stochastic noise Brownian motion w_t starting from $w_0 = 0$. The system state $x \in \mathbb{R}^n$ evolves according to the following stochastic differential equation (SDE)

$$dx = (f(x) + g(x)u)dt + \sigma dw, \quad (1)$$

where $u \in \mathbb{R}^m$ is the control input, σ is the magnitude of the noise. We assume that function f and g are parameterized by some parameter λ . Safety of the system is defined as the state staying within a safe set \mathcal{C} , which is the super-level set of a function $\phi(x) : \mathbb{R}^n \rightarrow \mathbb{R}$, *i.e.*,

$$\mathcal{C} = \{x \mid \phi(x) \geq 0\}. \quad (2)$$

This definition of safety can characterize a large variety of practical safety requirements (Ames et al., 2019). For the stochastic system (1), since safety can only be guaranteed in the sense of probability, we consider the long-term safety probability F_s and recovery probability F_r of the system defined as below.

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

$$F_s(x, T) = \mathbb{P}(x_t \in \mathcal{C}, \forall t \in [0, T] \mid x_0 = x). \quad (3)$$

Definition 2 (Recovery probability) *Starting from initial state $x_0 = x \notin \mathcal{C}$, the recovery probability F_r of system (1) for outlook time horizon T is defined as the probability of state get back to the safe set during the time interval $[0, T]$.*

$$F_r(x, T) = \mathbb{P}(\exists t \in [0, T], x_t \in \mathcal{C} \mid x_0 = x). \quad (4)$$

We point out that both safety probability and recovery probability are specific realizations of risk probability, depending on whether the initial state is safe or not, and whether the safe or unsafe events are of interest in the system (they are complementary). In the rest of the paper, we will denote risk probability as F , and it refers to safety probability or recovery probability depending on the initial state of the system is within the safe set or not.

The value of the risk probability F over the state space is crucial for safe control of these systems, but in practice it is hard to obtain all risk probability information accurately, *e.g.*, the failure probability of manufacturing robot arm is very small thus hard to estimate (Lasota et al., 2014). In this paper, the goal is to accurately estimate the risk probability and its gradient over the whole state space, and to achieve adaptation in changing system dynamics. Specifically, we want to find

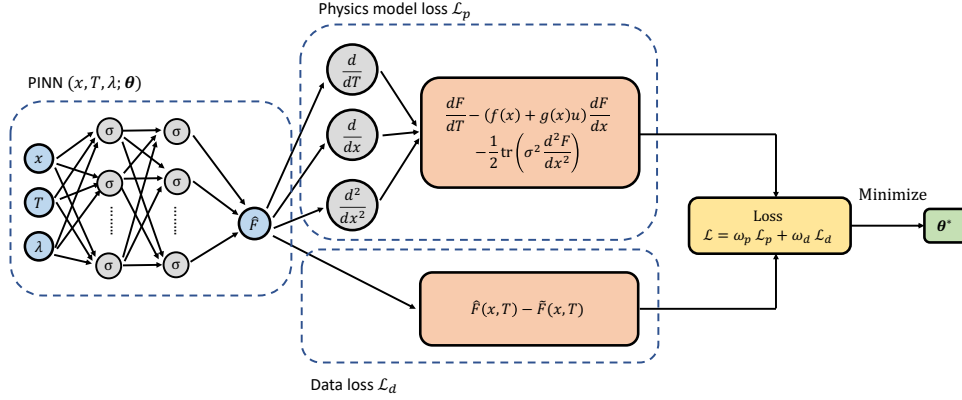


Figure 2: The training scheme of the physics-informed neural network (PINN) in PIPE. The PINN takes state x , time horizon T and the system parameter λ as input, and outputs the prediction of risk probability \hat{F} . The PINN is parameterized by θ . The loss function is the weighted sum of physics model loss \mathcal{L}_p and data loss \mathcal{L}_d . The physics model loss calculates the satisfaction of the governing PDE on sampled data points in the state-time space. The partial derivatives and Hessian of the risk probability come naturally from the automatic differentiation in deep learning frameworks such as PyTorch (Paszke et al., 2019) and TensorFlow (Abadi et al., 2016). At the end, we use an optimizer to minimize the loss function and find the optimal parameters in the PINN.

the mapping between the state-time pair (x, T) and the risk probability F over the entire state-time space $\mathbb{R}^n \times \mathbb{R}$ for any system parameter λ , and to estimate the gradient of risk probability F with regard to the state x . We split the problem into four tasks: generalization to unseen regions in the state-time space, efficient estimation with fixed number of sample data, adaptation on changing system parameters, and accurate estimation of probability gradients. Detailed settings can be found in experiment section of this paper. For all tasks, we will visualize the prediction results for qualitative information and calculate the average absolute error compared to the ground truth for quantitative analysis.

4. Proposed Method

We propose a model-based data-driven approach: PIPE, to address the tasks listed in the previous section. According to (Chern et al., 2021), the risk probability of system (1) is characterized by the solution of a PDE.

Theorem 3 *For a stochastic control affine system in the form of (1), the risk probability F (for both safety probability F_s defined in Definition 1 and recovery probability F_r defined in Definition 2) is characterized by the following convection diffusion equation.*

$$W(x, T) = \frac{\partial F}{\partial t}(x, T) - (f(x) + g(x)u) \frac{\partial F}{\partial x}(x, T) - \frac{1}{2} \text{tr} \left(\sigma^2 \frac{\partial^2 F}{\partial x^2}(x, T) \right) = 0, \quad (5)$$

with initial condition $F(x, 0) = \mathbb{1}(x \in \mathcal{C})$. For recovery probability, the boundary condition is $F(x, T) = 1$, $x \in \partial \mathcal{C}$. For safety probability, the boundary condition is $F(x, T) = 0$, $x \notin \mathcal{C}$.

Theorem 3 states that the risk probability of a control system can be analytically expressed as a PDE. The PDE is constituted of a convection term and a diffusion term. The convection term characterizes how the risk probability changes as the system evolves under its deterministic part of the dynamics. The diffusion term characterizes the effect of stochastic noises on the risk probability value as the noise term in the dynamics diffuses out with time. The initial condition says, when the outlook time T is 0, the risk probability value is the indicator function of whether the state is within the safe set. The boundary condition says that on the boundary of safe set \mathcal{C} , the risk probability F is 0 if we consider safety probability, and is 1 if we consider recovery probability. We use $W(x, T)$ to denote the function value that the PDE risk probability F should satisfy, to better define the loss function in the learning framework later.

While the PDE provides a way to get the actual safety probability of the system, to solve a PDE is not easy in general, especially when the coefficient of the PDE is time varying, as in the case of (5). Despite the low sample efficiency, MC methods provide another way to solve this problem. Assume the dynamics of the system is given, one can simulation the system for an initial condition multiple times to get an empirical estimate of the safety probability, by calculating the ratio of safe trajectories over all trajectories. However, MC requires huge number of trajectories to get accurate estimation, and the evaluation of safety probability can only be conducted at a single point at a time.

To leverage the advantages of PDE and MC, and to overcome their drawbacks, we propose to use physics-informed neural networks (PINNs) to learn the mapping from state-time horizon pair (x, T) to the risk probability value F . Fig. 2 shows the architecture of the PINN. The PINN takes the state-time pair (x, T) and the system parameter λ as the input, and outputs the safety probability prediction \hat{F} , the state and time derivatives $\frac{\partial \hat{F}}{\partial x}$ and $\frac{\partial \hat{F}}{\partial T}$, and the Hessian $\frac{\partial^2 \hat{F}}{\partial x^2}$. Unlike standard PINN, we add the system parameter λ as an input to achieve adaptations on varying system parameters. Assume the PINN is parameterized by θ . The loss function is defined as

$$\mathcal{L}(\theta) = \omega_p \mathcal{L}_p(\theta) + \omega_d \mathcal{L}_d(\theta), \quad (6)$$

where

$$\begin{aligned} \mathcal{L}_p(\theta) &= \frac{1}{\mathcal{P}} \sum_{(x,T) \in \mathcal{P}} \|W(x, T)\|_2^2, \\ \mathcal{L}_d(\theta) &= \frac{1}{\mathcal{D}} \sum_{(x,T) \in \mathcal{D}} \|\hat{F}(x, T) - \tilde{F}(x, T)\|_2^2. \end{aligned} \quad (7)$$

Here, \mathcal{P} and \mathcal{D} are the training point sets for physics model and external data, respectively. The loss function \mathcal{L} is constituted by two parts, 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 $W(x, T)$, 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 ω_p and ω_d .

The resulting PIPE framework combines MC data and the governing PDE into a PINN, to directly learn a mapping from the state-time pair to the risk probability. The advantages of the PIPE framework include fast inference at test time, accurate estimation, and ability to generalize from the combination of data and model.

5. Experiments

We conduct four experiments to illustrate the efficacy of the proposed method. The system dynamics of interest is (1) with $x \in \mathbb{R}$, $f(x) = \lambda dt$, $g(x) = 0$ and $\sigma = 1$. The system dynamics become

$$dx = \lambda dt + dw. \quad (8)$$

The safe set is defined as (2) with $\phi(x) = x - 2$. The state-time region of interest is $\Omega \times \tau = [-10, 2] \times [0, 10]$. For risk probability, we consider the recovery probability of the system from initial state $x_0 \notin \mathcal{C}$ outside the safe set. Specifically, the risk probability F is characterized by the solution of the following convection diffusion equation

$$\frac{\partial F}{\partial T}(x, T) = \lambda \frac{\partial F}{\partial x}(x, T) + \frac{1}{2} \text{tr} \left(\frac{\partial^2 F}{\partial x^2}(x, T) \right), \quad (9)$$

with initial condition $F(x, 0) = \mathbb{1}(x \geq 2)$ and boundary condition $F(2, T) = 1$. We choose this system because we have the analytical solution of (9) as ground truth for comparison, as given by

$$F(x, T) = \frac{(2 - x) \exp \left\{ \frac{-((2-x)-\lambda T)^2}{2T} \right\}}{\sqrt{2\pi T^3}}. \quad (10)$$

The empirical data of the risk probability is acquired by running MC with the system dynamics (1) with initial state $x = x_0$ multiple times, and calculate the number of trajectories that the state recovers to safe set during the time horizon $[0, T]$ over the full trajectory number, *i.e.*,

$$\tilde{F}(x, T) = \mathbb{P}(\exists t \in [0, T], x_t \in \mathcal{C} \mid x_0 = x) = \frac{N_{\text{recovery}}}{N}, \quad (11)$$

where N is the number of sample trajectories, and is a tunable parameter that affects the accuracy of the estimated risk probability. Specifically, larger N gives more accurate estimation.

In all experiments, we use PINN with 3 hidden layers and 32 neurons per layer. The activation function is chosen as hyperbolic tangent function (tanh). We use Adam optimizer (Kingma and Ba, 2014) for training with initial learning rate set as 0.001. The PINN parameters θ is initialized via Glorot uniform initialization. The weights in the loss function (6) are set to be $\omega_p = \omega_d = 1$. We train the PINN for 60000 epochs in all experiments. The simulation is constructed based on the DeepXDE framework (Lu et al., 2021). Details about the experiments, simulation results on other systems, and applications to stochastic safe control can be found at the project page: <https://github.com/jacobwang925/PIPE-L4DC>.

5.1. Generalization to unseen regions

In this experiment, we want to test the generalization ability of PIPE to unseen regions of the state-time space. We consider system (8) with $\lambda = 1$. We train the PINN with data only on the sub-region of the state-time space $\Omega \times \tau = [-10, -2] \times [0, 10]$, but test the trained PINN on the full state-time region $\Omega \times \tau = [-10, 2] \times [0, 10]$. The training data is acquired through MC with sample trajectories number $N = 1000$, and is down-sampled to $dx = 0.4$ and $dt = 0.5$. For comparison, we use thin plate spline (TPS) fitting on the training data to infer the risk probability on the whole state space. We also examined other fitting methods such as cubic spline, polynomial fitting, *etc.*




Figure 3: (I) Settings and results of the risk probability generalization task. PIPE and TPS fitting are compared. The average absolute error of prediction is 9.2×10^{-2} for TPS, and 0.3×10^{-2} for PIPE. (II) Gradient of the risk probability prediction of PIPE and MC. The average absolute error of gradient prediction is 2.78×10^{-2} for MC, and 0.06×10^{-2} for PIPE.

but TPS performs the best over all fitting strategies. Fig. 3 (I) visualizes the training data samples and shows the results. The spline fitting does not include any physical model that the system should satisfy, thus fails to generalize to unseen regions in the state space. On contrary, PIPE can infer the risk probability value very accurately on the whole state space, due to its combination of data and physic model.

5.2. Efficient estimation of risk probability

In this experiment, we show that PIPE will give more efficient estimations of risk probability in terms of accuracy and sample number, compared to MC and its variants. We consider system (8) with $\lambda = 1$. The training data is sampled on the state-time space $\Omega \times \tau = [-10, 2] \times [0, 10]$ with $dx = 0.2$ and $dt = 0.1$. We compare the risk probability estimation error of PIPE and MC, on two regions in the state-time space:

1. Normal event region: $\Omega \times \tau = [-6, -2] \times [4, 6]$ where the average probability is 0.412.
2. Rare event region: $\Omega \times \tau = [-2, 0] \times [8, 10]$ where the average probability is 0.985.

For fairer comparison, we use a uniform filter of kernel size 3 on the MC data to smooth out the noise, as the main cause of inaccuracy of MC estimation is sampling noise. Fig. 4 shows the percentage errors of risk probability inference under different MC sample numbers N . As the sample number goes up, prediction errors for all three approaches decrease. The denoised MC has lower error compared to standard MC as a result of denoising, and their errors tend to converge




Figure 4: Percentage error of risk probability estimation for different MC sample numbers for rare events and normal events. PIPE, MC and denoised MC with uniform kernel filtering are compared. Both error and sample number are in log scale.

since the sampling noise contributes less to the error as the sample number increases. On both rare events and normal events, PIPE yields more accurate estimation than MC and denoised MC across all sample numbers. This indicates that PIPE has better sample efficiency than MC and its variants, in the sense that it requires less sample data to achieve same prediction accuracy. This desired feature of PIPE is due to the fact that PIPE incorporates model knowledge into the MC data to further enhance its accuracy, by taking the physics-informed neighboring relationships of the data into consideration.

5.3. Adaptation on changing system parameters

In this experiment, we show that the proposed PIPE framework will allow generalization to uncertain parameters of the system. We consider system (8) with varying $\lambda \in [0, 2]$. We use MC data with sample number $N = 10000$ for a fixed set of $\lambda_{\text{train}} = [0.1, 0.5, 0.8, 1]$ for training, and test PIPE after training on $\lambda_{\text{test}} = [0.3, 0.7, 1.2, 1.5, 2]$. We only present $\lambda_{\text{test}} = 1.5$ due to space limit, similar results on different λ_{test} can be found at the project webpage. Fig. 5 shows the results. We can see that PIPE is able to predict risk probability for systems with unseen and even out of distribution parameters during training. In the prediction error plot, the only region that has larger prediction error is at $T = 0$ and $x \in \partial\mathcal{C}$ on the boundary of the safe set. This is because the risk probability at this point is not well defined (it can be either 0 or 1), and this point will not be considered in a control scenario as we are always interested in long-term safety where $T \gg 0$. This adaptation feature of the PIPE framework indicates its potential use on stochastic safe control with uncertain system parameters, and it also opens the door for physics-informed learning on a family of PDEs. In general, PDEs with different parameters can have qualitatively different behaviors, so is hard to generalize. The control theory model allows us to have a sense when the PDEs are qualitatively similar with different parameters, and thus allows generalization within the qualitatively similar cases.

5.4. Estimating the gradient of risk probability

In this experiment, we show that PIPE is able to generate accurate gradient predictions of risk probabilities. We consider system (8) with $\lambda = 1$. Similar to the generalization task, we train the PINN with MC data of $N = 1000$ on the sub-region $\Omega \times \tau = [-10, -2] \times [0, 10]$ and test the trained PINN on the full state-time region $\Omega \times \tau = [-10, 2] \times [0, 10]$. We then take finite difference of the risk probability with regard to the state x to calculate its gradient, for ground truth F , MC estimation \hat{F} and PIPE prediction \hat{F} . Fig. 3 (II) shows the results. It can be seen that PIPE gives much more accurate estimation of risk probability gradient than MC, and this is due to the fact




Figure 5: Risk probability prediction of PIPE on unseen system parameters. The average absolute error of the prediction is 0.70×10^{-2} .

that PIPE incorporates physics model information inside the training process. It is also interesting that PIPE does not use any governing laws of the risk probability gradient during training, and by considering the risk probability PDE alone, it can provide very accurate estimation of the gradient. The results indicates that PIPE can enable the usage a lot of first- and higher-order stochastic safe control methods online, by providing accurate and fast estimation of the risk probability gradients.

6. Discussion

The results indicate several potential applications of the proposed method. Since PIPE generalizes to unseen region, we can use it to accurately estimate safety probability value for regions where it is hard to get data (*e.g.*, dangerous working zone for robotic systems, rare unsafe events in industrial manufacturing). With the adaptation ability of the proposed framework, we are able to deal with systems with uncertain parameters, and produce fast estimation and response when the system parameters are changing over time. Also, a lot of gradient-based methods can be used for stochastic safe control combining with existing optimizers (*e.g.*, stochastic gradient descents), since the proposed PIPE is able to accurately estimate the risk probability gradients.

Furthermore, the proposed PIPE framework is generic, and can be potentially applied to other different scenarios beyond safe control. The adaptation ability over the changing parameters in the system indicates some underlying relationships of the PDE solution with different coefficients but under the same structure, which opens up exciting future research directions on theoretical analysis of the learning behaviours of PINN.

7. Conclusion

In this paper, we proposed a generalizable physics-informed learning framework: PIPE, to estimate the risk probability in stochastic safe control systems. The proposed PIPE framework combines data from Monte Carlo and the underlying governing PDE of the risk probability, to accurately learn the risk probability as well as its gradient. PIPE has better sample efficiencies compared to MC, and is able to generalize to unseen regions in the state space beyond training. The resulting PIPE framework is also robust to uncertain parameters in the system dynamics, and can infer risk probability values of a class of systems with training data only from a fixed number of systems. The proposed PIPE framework provides key foundations for first- and higher-order methods for stochastic safe control, and opens the door for robust physics-informed learning for generic PDEs. Future work includes applications to high-dimensional and real-world 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.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- Alessandro Abate, Maria Prandini, John Lygeros, and Shankar Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.
- Mohamadreza Ahmadi, Xiaobin Xiong, and Aaron D Ames. Risk-averse control via cvar barrier functions: Application to bipedal robot locomotion. *IEEE Control Systems Letters*, 6:878–883, 2021.
- Aaron D Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control barrier functions: Theory and applications. In *2019 18th European control conference (ECC)*, pages 3420–3431. IEEE, 2019.
- Siu-Kui Au and James L Beck. Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic engineering mechanics*, 16(4):263–277, 2001.
- Zdravko I Botev, Pierre L’Ecuyer, and Bruno Tuffin. Markov chain importance sampling with applications to rare event probability estimation. *Statistics and Computing*, 23(2):271–285, 2013.
- Shengze Cai, Zhiping Mao, Zhicheng Wang, Minglang Yin, and George Em Karniadakis. Physics-informed neural networks (pinns) for fluid mechanics: A review. *Acta Mechanica Sinica*, pages 1–12, 2022.
- Frédéric Cérou, Pierre Del Moral, Teddy Furon, and Arnaud Guyader. Sequential monte carlo for rare event estimation. *Statistics and computing*, 22(3):795–808, 2012.
- Margaret P Chapman, Jonathan Lacotte, Aviv Tamar, Donggun Lee, Kevin M Smith, Victoria Cheng, Jaime F Fisac, Susmit Jha, Marco Pavone, and Claire J Tomlin. A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems. In *2019 American Control Conference (ACC)*, pages 2958–2963. IEEE, 2019.
- Albert Chern, Xiang Wang, Abhiram Iyer, and Yorie Nakahira. Safe control in the presence of stochastic uncertainties. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 6640–6645. IEEE, 2021.
- William S Cleveland. Lowess: A program for smoothing scatterplots by robust locally weighted regression. *American Statistician*, 35(1):54, 1981.

- Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Hans Janssen. Monte-carlo based uncertainty analysis: Sampling efficiency and sampling convergence. *Reliability Engineering & System Safety*, 109:123–132, 2013.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Przemyslaw A Lasota, Gregory F Rossano, and Julie A Shah. Toward safe close-proximity human-robot interaction with standard industrial robots. In *2014 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 339–344. IEEE, 2014.
- Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. Deepxde: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021.
- George S Misyris, Andreas Venzke, and Spyros Chatzivasileiadis. Physics-informed neural networks for power systems. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.
- Yashwanth Kumar Nakka, Anqi Liu, Guanya Shi, Anima Anandkumar, Yisong Yue, and Soon-Jo Chung. Chance-constrained trajectory optimization for safe exploration and learning of nonlinear systems. *IEEE Robotics and Automation Letters*, 6(2):389–396, 2020.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- Marcus Pereira, Ziyi Wang, Ioannis Exarchos, and Evangelos Theodorou. Safe optimal control using stochastic barrier functions and deep forward-backward sdes. In *Conference on Robot Learning*, pages 1783–1801. PMLR, 2021.
- Samuel Pfrommer, Tanmay Gautam, Alec Zhou, and Somayeh Sojoudi. Safe reinforcement learning with chance-constrained model predictive control. In *Learning for Dynamics and Control Conference*, pages 291–303. PMLR, 2022.
- Maria Prandini and Jianghai Hu. Application of reachability analysis for stochastic hybrid systems to aircraft conflict prediction. In *2008 47th IEEE Conference on Decision and Control*, pages 4036–4041. IEEE, 2008.
- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. 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, 2019.
- Gerardo Rubino and Bruno Tuffin. *Rare event simulation using Monte Carlo methods*. John Wiley & Sons, 2009.
- Francisco Sahli Costabal, Yibo Yang, Paris Perdikaris, Daniel E Hurtado, and Ellen Kuhl. Physics-informed neural networks for cardiac activation mapping. *Frontiers in Physics*, 8:42, 2020.

- Samantha Samuelson and Insoon Yang. Safety-aware optimal control of stochastic systems using conditional value-at-risk. In *2018 Annual American Control Conference (ACC)*, pages 6285–6290. IEEE, 2018.
- Cesar Santoyo, Maxence Dutreix, and Samuel Coogan. A barrier function approach to finite-time stochastic system verification and control. *Automatica*, 125:109439, 2021.
- Andrew Singletary, Mohamadreza Ahmadi, and Aaron D Ames. Safe control for nonlinear systems with stochastic uncertainty via risk control barrier functions. *arXiv preprint arXiv:2203.15892*, 2022.
- Zhuoyuan Wang, Haoming Jing, Christian Kurniawan, Albert Chern, and Yorie Nakahira. Myopically verifiable probabilistic certificate for long-term safety. *arXiv preprint arXiv:2110.13380*, 2021.
- Hai Zhu and Javier Alonso-Mora. Chance-constrained collision avoidance for mavs in dynamic environments. *IEEE Robotics and Automation Letters*, 4(2):776–783, 2019.
- Konstantin Zuev. Subset simulation method for rare event estimation: an introduction. *arXiv preprint arXiv:1505.03506*, 2015.

Appendix A. Experiment Details

In this section, we provide details for the four experiments presented in the paper.

A.1. Generalization to unseen regions

In the generalization task in section 5.1, we use a down-sampled sub-region of the system to train the proposed PIPE framework, and test the prediction results on the full state-time space. We showed that PIPE is able to give accurate inference on the entire state space, while standard fitting strategies cannot make accurate predictions. In the paper, we only show the fitting results of thin plate spline interpolation. Here, we show the results of all fitting strategies we tested for this generalization tasks. The fitting strategies are

1. Polynomial fitting of 5 degrees for both state x and time T axes. The fitting sum of squares error (SSE) on the training data is 0.1803.
2. Lowess: locally weighted scatterplot smoothing (Cleveland, 1981). The training SSE is 0.0205.
3. Cubic spline interpolation. The training SSE is 0.
4. Biharmonic spline interpolation. The training SSE is 2.52×10^{-27} .
5. TPS: thin plate spline interpolation. The training SSE is 1.64×10^{-26} .

All fittings are conducted via the MATLAB Curve Fitting Toolbox. Fig. 6 visualizes the fitting results on the full state space. Polynomial fitting performs undesirably because the polynomial functions cannot capture the risk probability geometry well. Lowess fitting also fails at inference since it does not have any model information of the data. Given the risk probability data, cubic spline cannot extrapolate outside the training region, and we use 0 value to fill in the unseen region where it yields NAN for prediction. Biharmonic and TPS give similar results as they are both spline interpolation methods. None of these fitting methods can accurately predict the risk probability in unseen regions, because they purely rely on training data and do not incorporate any model information of the risk probability for prediction.

We also compare the prediction results for different network architectures in the proposed PIPE framework, to examine the effect of network architectures on the risk probability prediction performance. The network settings we consider are different hidden layer numbers (1-4) and different numbers of neurons in one hidden layer (16, 32, 64). We use 3 hidden layer, 32 neuron per layer as baseline (the one used in the paper). Table 1 and Table 2 report the averaged absolute error of the predictions for different layer numbers and neuron numbers per layer, respectively. We trained the neural networks for 10 times with random initialization to report the standard deviation. We can see that as the number of layer increases, the prediction error of the risk probability drops, but in a relatively graceful manner. The prediction error for a single layer PIPE is already very small, which indicates that the proposed PIPE framework is very efficient in terms of computation and storage. The prediction accuracy tends to saturate while the hidden layer number reaches 3, as there is no obvious improvement when we increase the layer number from 3 to 4. This means for the specific task we consider, a 3-layer neural net has enough representation. Under the same layer number, as the neuron number per layer increases, the risk probability prediction error decreases. This indicates

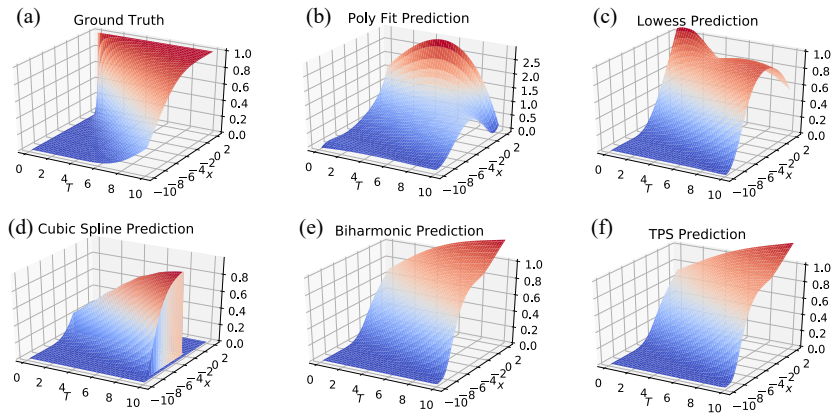


Figure 6: Results of different fitting strategies on the risk probability generalization task.

that with larger number of neuron in each layer (*i.e.*, wider neural networks), the neural network can better capture the mapping between state-time pair and the risk probability. However, the training time increases significantly for PIPEs with more neurons per layer (152s for 16 neurons and 971s for 64 neurons), and the gain in prediction accuracy becomes marginal compared to the amount of additional computation involved. We suggest to use a moderate number of neurons per layer to achieve desirable trade-offs between computation and accuracy.

# Hidden Layer	1	2	3	4
Prediction Error ($\times 10^{-3}$)	4.773 \pm 0.564	2.717 \pm 0.241	2.819 \pm 0.619	2.778 \pm 0.523

Table 1: Risk probability prediction error of PIPE for different numbers of hidden layers.

# Neurons	16	32	64
Prediction Error ($\times 10^{-3}$)	2.743 \pm 0.313	2.931 \pm 0.865	2.599 \pm 0.351

Table 2: Risk probability prediction error of PIPE for different neuron numbers per layer.

A.2. Efficient estimation of risk probability

In the efficient estimation task in section 5.2, we showed that PIPE will give better sample efficiency in risk probability prediction, in the sense that it achieves the same prediction accuracy with less sample numbers. Here, we visualize the prediction errors of Monte Carlo (MC) and the proposed PIPE framework to better show the results. Fig. 7 shows the prediction error comparison plots for MC and PIPE with different sample numbers N . As the sample number increases, the errors for both MC and PIPE decrease because of better resolution of the data. PIPE gives more accurate predictions than MC across all sample numbers, since it combines data and physical model of the system together. From Fig. 7 we can see that, PIPE indeed provides smoother and more accurate estimation of the risk probability. The visualization results further validate the efficacy of the proposed PIPE framework.

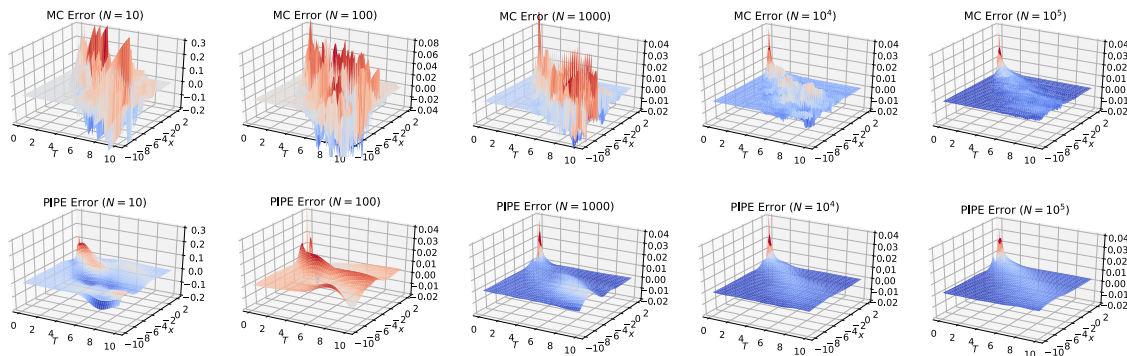


Figure 7: Prediction errors for Monte Carlo (left) and PIPE (right) with different sample numbers.

# Hidden Layer	1	2	3	4
Prediction Error ($\times 10^{-4}$)	14.594 ± 2.109	7.302 ± 0.819	6.890 ± 0.613	6.625 ± 0.574

Table 3: Risk probability gradient prediction error of PIPE for different numbers of hidden layers.

A.3. Adaptation on changing system parameters

For the adaptation task described in section 5.3, we trained PIPE with system data of parameters $\lambda_{\text{train}} = [0.1, 0.5, 0.8, 1]$ and tested over a range of unseen parameters over the interval $\lambda = [0, 2]$. Here, we show additional results on parameters $\lambda_{\text{test}} = [0.3, 0.7, 1.2, 2]$ to further illustrate the adaptation ability of PIPE. Fig. 8 shows the results. It can be seen that PIPE is able to predict the risk probability accurately on both system parameters with very low error over the entire state-time space. This result indicates that PIPE has solid adaptation ability on uncertain parameters, and can be used for stochastic safe control with adaptation requirements.

A.4. Estimating the gradient of risk probability

For the gradient estimation task in section 5.4, we presented that PIPE is able to predict the risk probability gradients accurately by taking finite difference on the predicted risk probabilities. This result shows that PIPE can be used for first- and higher-order methods for safe control, by providing accurate gradient estimations in a real-time fashion. Similar to the generalization task, here we report the gradient prediction errors with different network architectures in PIPE, to examine the effect of network architectures on the gradient estimation performance. Table 3 and Table 4 show the averaged absolute error of gradient predictions for different layer numbers and neuron numbers per layer. We trained the neural networks for 10 times with random initialization to report the standard deviation. We can see that as the number of hidden layer increases, the gradient prediction error keeps dropping, and tends to saturate after 3 layers. With the increasing neuron numbers per layer, the gradient prediction error decreases in a graceful manner. Similar to the generalization task, even though a larger networks with more hidden layers and more neurons per layer can give more accurate estimation of the gradient, the computation scales poorly compared to the accuracy

# Neurons	16	32	64
Prediction Error ($\times 10^{-4}$)	7.049 ± 0.767	6.890 ± 0.613	6.458 ± 0.794

Table 4: Risk probability gradient prediction error of PIPE for different neuron numbers per layer.




Figure 8: Risk probability prediction results on systems with unseen parameters $\lambda_{\text{test}} = [0.3, 0.7, 1.2, 2]$.

gain. Based on those results, we suggest to use moderate numbers of layers and neurons per layer to acquire desirable gradient prediction with less computation time.

Appendix B. Additional Results

In this section we provide additional experiment results on safe control with risk estimation from the proposed PIPE framework.

Appendix C. Risk estimation of inverted pendulum on a cart system

We consider the inverted pendulum on a cart system, with dynamics being

$$\frac{d\mathbf{x}}{dt} = f(\mathbf{x}) + g(\mathbf{x})u + \sigma \tilde{I} dW_t, \quad (12)$$

where $\mathbf{x} = [x, \dot{x}, \theta, \dot{\theta}]^\top$ is the state of the system and $u \in \mathbb{R}$ is the control, with x and \dot{x} being the position and velocity of the cart, θ and $\dot{\theta}$ being the angle and angular velocity of the pendulum. We use \mathbf{x} to denote the state of the system to distinguish from cart's position x . And we have

$$f(\mathbf{x}) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & m + M & 0 & ml \cos \theta \\ 0 & 0 & 1 & 0 \\ 0 & ml \cos \theta & 0 & ml^2 \end{bmatrix}^{-1} \begin{bmatrix} \dot{x} \\ ml\dot{\theta}^2 \sin \theta - b_x \dot{x} \\ \dot{\theta} \\ mgl \sin \theta - b_\theta \dot{\theta} \end{bmatrix}, \quad (13)$$

$$g(\mathbf{x}) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & m + M & 0 & ml \cos \theta \\ 0 & 0 & 1 & 0 \\ 0 & ml \cos \theta & 0 & ml^2 \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad (14)$$




Figure 9: Inverted pendulum on a cart.

where m and M are the mass of the pendulum and the cart, g is the acceleration of gravity, l is the length of the pendulum, and b_x and b_θ are constants. The last term in (12) is the additive noise, where W_t is 4-dimensional Brownian motion with $W_0 = \mathbf{0}$, σ is the magnitude of the noise, and

$$\tilde{I} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (15)$$

Fig. 9 visualizes the system.

The safe set is defined in (2) with barrier function being $\phi(\mathbf{x}) = 1 - (\frac{x_3}{\pi/3})^2 = 1 - (\frac{\theta}{\pi/3})^2$. Essentially the system is safe when the angle of the pendulum is within $[-\pi/3, \pi/3]$. We consider spatial-temporal space $\Omega \times \tau = [[-10, 10] \times [-\pi/3, \pi/3] \times [-\pi, \pi]] \times [0, 1]$. We collect training and testing data on $\Omega \times \tau$ with grid size $N_{\Omega\text{-train}} = 13$ and $N_{\tau\text{-train}} = 10$ for training and $N_{\Omega\text{-test}} = 25$ and $N_{\tau\text{-test}} = 10$ for testing. The nominal controller we choose is a proportional controller $N(\mathbf{x}) = -K\mathbf{x}$ with $K = [0, -0.9148, -22.1636, -14.3992]^\top$. The sample number for MC simulation is set to be $N = 1000$ for both training and testing. Table 5 lists the parameters used in the simulation.

We train PIPE with the same configuration listed in section 5. According to Theorem 3, the PDE that characterizes the safety probability of the pendulum system is

$$\frac{\partial F}{\partial T}(\mathbf{x}, T) = (f(\mathbf{x}) - g(\mathbf{x})K\mathbf{x}) \frac{\partial F}{\partial \mathbf{x}}(\mathbf{x}, T) + \frac{1}{2}\sigma^2 \tilde{I} \text{tr} \left(\frac{\partial^2 F}{\partial \mathbf{x}^2}(\mathbf{x}, T) \right), \quad (16)$$

which is a high dimensional and highly nonlinear PDE that can not be solved effectively using standard solvers. Here we can see the advantage of combining data and model and using a learning-based framework to estimate the safety probability. Fig. 10, Fig. 11 and Fig. 12 show the results of the PIPE predictions. We can see that despite the rather high dimension and the nonlinear dynamics of the pendulum system, PIPE is able to predict the safety probability of the system with high accuracy. Besides, since PIPE takes the model knowledge into training loss, the resulting safety probability prediction is smoother thus more reliable than pure MC estimations.

Parameters	Values
M	1
m	0.1
g	9.8
l	0.5
b_x	0.05
b_θ	0.1

Table 5: Parameters used in the inverted pendulum simulation.




Figure 10: Safety probability from MC simulation and PIPE prediction, and their difference. Results on outlook time horizon $T = 0.6$, initial velocity $v = 0$. The x-axis shows the initial angle, and y-axis shows the initial angular velocity.




Figure 11: Safety probability from MC simulation and PIPE prediction, and their difference. Results on outlook time horizon $T = 0.3$. The 5x5 plots show results on 25 different initial velocity uniformly sampled in $[-10, 10]$. The x-axis and y-axis (omitted) are the initial angle and the initial angular velocity as in Fig. 10. One can see that the safety probability shift as the velocity changes, and the safety probability is symmetric with regard to the origin $v = 0$ due to the symmetry of the dynamics.




Figure 12: Safety probability from MC simulation and PIPE prediction, and their difference. Results on outlook time horizon $T = 1$. The 5x5 plots show results on 25 different initial velocity uniformly sampled in $[-10, 10]$.

C.1. Safe control with PIPE

We consider the control affine system (1) with $f(x) \equiv Ax = 2x$, $g(x) \equiv 1$, $\sigma(x) \equiv 2$. The safe set is defined in (2) and the barrier function is chosen to be $\phi(x) := x - 1$. The safety specification is given as the forward invariance condition. The nominal controller is a proportional controller $N(x) = -Kx$ with $K = 2.5$. The closed-loop system with this controller has an equilibrium at $x = 0$ and tends to move into the unsafe set in the state space. We run simulations with $dt = 0.1$ for all controllers. The initial state is set to be $x_0 = 3$. For this system, the safety probability satisfies the following PDE

$$\frac{\partial F}{\partial t}(x, T) = -0.5x \frac{\partial F}{\partial x}(x, T) + 2 \operatorname{tr} \left(\frac{\partial^2 F}{\partial x^2}(x, T) \right), \quad (17)$$

with initial and boundary conditions

$$\begin{aligned} F(x, t) &= 0, \quad x \leq 1, \\ F(x, 0) &= \mathbb{1}(x \geq 1). \end{aligned} \quad (18)$$

We first estimate the safety probability $F(x, T)$ of the system via PIPE. The training data $\tilde{F}(x, T)$ is acquired by running MC on the system dynamics for given initial state x_0 and nominal control N . Specifically,

$$\tilde{F}(x, T) = \mathbb{P}(\forall t \in [0, T], x_t \in \mathcal{C} \mid x_0 = x) = \frac{N_{\text{safe}}}{N}, \quad (19)$$

with $N = 100$ be the number of sample trajectories. The training data is sampled on the state-time region $\Omega \times \tau = [1, 10] \times [0, 10]$ with grid size $dx = 0.5$ and $dt = 0.5$. We train PIPE with the same configuration as listed in section 5. We test the estimated safety probability and its gradient on the full state space $\Omega \times \tau$ with $dx = 0.1$ and $dt = 0.1$. Fig. 13 shows the results. It can be seen that the PIPE estimate is very close to the Monte Carlo samples, which validates the efficacy of the framework. Furthermore, the PIPE estimation has smoother gradients, due to the fact that it leverages model information along with the data.

We then show the results of using such estimated safety probability for control. Fig. 14 shows the results. For the baseline stochastic safe controller for comparison, refer to (Wang et al., 2021) for details. We can see that with the control method in (Wang et al., 2021) together with the PIPE estimation, long-term safety can be ensured.

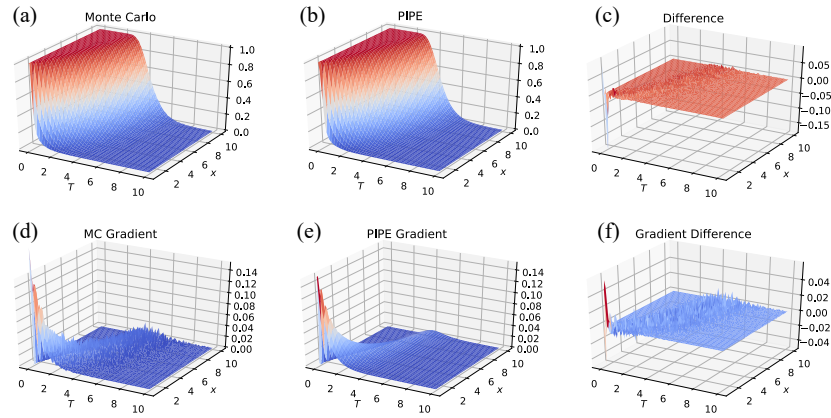


Figure 13: Safety probability and its gradient of Monte Carlo samples and PIPE estimation.




Figure 14: Safe control with probability estimation from PIPE, compared with other baselines.