

PINF: CONTINUOUS NORMALIZING FLOWS FOR PHYSICS-CONSTRAINED DEEP LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

The normalization constraint on probability density poses a significant challenge for solving the Fokker-Planck equation. Normalizing Flow, an invertible generative model leverages the change of variables formula to ensure probability density conservation and enable the learning of complex data distributions. In this paper, we introduce Physics-Informed Normalizing Flows (PINF), a novel extension of continuous normalizing flows, incorporating diffusion through the method of characteristics. Our method, which is mesh-free and causality-free, can efficiently solve high dimensional time-dependent and steady-state Fokker-Planck equations.

1 INTRODUCTION

The Fokker-Planck (FP) equation (Risken, 1996) is a well-known partial differential equation that describes the evolution of a stochastic system’s probability density function (PDF) over time. Due to the high-dimensional variable and unbounded domain, traditional numerical methods, such as the finite difference methods (FDM) (Kumar & Narayanan, 2006), the finite element methods (FEM) (Deng, 2009), and the path integral methods (Wehner & Wolfer, 1983) prove to be computationally daunting in tackling the FP equations. By contrast, deep learning algorithms (Sirignano & Spiliopoulos, 2018; Raissi et al., 2019; Weinan & Yu, 2018), without a specific network structure, violate the normalization constraint on PDF, resulting in reduced accuracy and prevalent errors.

Recently, researchers have recognized the potential of flow-based generative models in learning complicated probability distributions (Kingma & Dhariwal, 2018; Ho et al., 2019; Albergo et al., 2019), alongside the connection to optimal transport (OT) theory (Finlay et al., 2020; Yang & Karniadakis, 2020; Zhang et al., 2018). Tang et al. (2022) employed normalizing flows as alternative solutions and proposed the KRnet for solving the steady-state Fokker-Planck (SFP) equations. However, the discrete normalizing flow remains circumscribed to model a single target distribution at a time, typically the final or steady-state distribution, thus constraining its utility in the time-dependent Fokker-Planck (TFP) equations. Consequently, Feng et al. (2022) introduced the temporal normalizing flow (TNF) to estimate time-dependent distributions for TFP equations, albeit without encapsulating the inherent physical laws of FP equations in structure.

More naturally, we generalize the continuous normalizing flow (CNF) (Chen et al., 2018) with diffusion and propose a novel intelligent architecture: Physics-Informed Normalizing Flows (PINF) for solving FP equations. We encode the physical constraints into ordinary differential equations (ODEs) using the method of characteristics and train the model in a self-supervised way. Numerical experiments demonstrate both accuracy and efficiency in solving high-dimensional TFP and SFP equations, even without the need for meticulous hyperparameter tuning.

2 PROBLEM DEFINITION

Here we first give a brief introduction to the FP equation. Consider the state variable $\mathbf{X}_t \in \mathbb{R}^d$ described by the stochastic differential equation (SDE) (Oksendal, 2013):

$$d\mathbf{X}_t = \mu(\mathbf{X}_t, t)dt + \sigma(\mathbf{X}_t, t)d\mathbf{W}_t, \quad (1)$$

where the drift coefficient $\mu(\mathbf{X}_t, t) \in \mathbb{R}^d$ is a vector field, $\sigma(\mathbf{X}_t, t) \in \mathbb{R}^{d \times M}$ is a matrix-valued function and \mathbf{W}_t is an M -dimensional standard Wiener process. The probability density function

$p(\mathbf{x}, t)$ of \mathbf{X}_t satisfies the following Fokker-Planck equation:

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla \cdot [p(\mathbf{x}, t)\mu(\mathbf{x}, t)] + \nabla \cdot [\nabla \cdot (p(\mathbf{x}, t)\mathbf{D}(\mathbf{x}, t))], \quad \forall(\mathbf{x}, t) \in \mathbb{R}^d \times \mathbb{R}^+, \\ p(\mathbf{x}, 0) &= p_0(\mathbf{x}), \\ p(\mathbf{x}) &\rightarrow 0 \quad \text{as } \|\mathbf{x}\| \rightarrow \infty, \end{aligned} \quad (2)$$

where $(\mathbf{x}, t) \in \mathbb{R}^{d+1}$ denote the spatial-temporal variables, $\mathbf{D}(\mathbf{x}, t) = \frac{1}{2}\sigma(\mathbf{x}, t)\sigma(\mathbf{x}, t)^T$ is the diffusion matrix, $p_0(\mathbf{x})$ is the initial PDF of \mathbf{X}_t , $p(\mathbf{x}, t)$ is defined with the unbounded boundary condition, ∇ is an operator for spatial variables and $\|\mathbf{x}\|$ indicates the ℓ_2 norm of \mathbf{x} .

The stationary solution of Eq.(2) means the invariant measure independent of time, satisfying

$$-\nabla \cdot [p(\mathbf{x}, t)\mu(\mathbf{x}, t)] + \nabla \cdot [\nabla \cdot (p(\mathbf{x}, t)\mathbf{D}(\mathbf{x}, t))] = 0, \quad \forall(\mathbf{x}, t) \in \mathbb{R}^d \times \mathbb{R}^+ \quad (3)$$

More specifically, i.e.,

$$-\sum_{i=1}^d \frac{\partial}{\partial x_i} [p(\mathbf{x}, t)\mu_i(\mathbf{x}, t)] + \sum_{i=1}^d \sum_{j=1}^d \frac{\partial^2}{\partial x_i \partial x_j} [p(\mathbf{x}, t)D_{ij}(\mathbf{x}, t)] = 0, \quad \forall(\mathbf{x}, t) \in \mathbb{R}^d \times \mathbb{R}^+ \quad (4)$$

For the physical background of the FP equation, there are some extra constraints on its solution $p(\mathbf{x}, t)$:

$$\begin{aligned} \int_{\mathbb{R}^d} p(\mathbf{x}, t) d\mathbf{x} &= 1, \quad \forall t \in \mathbb{R}^+ \\ p(\mathbf{x}, t) &\geq 0, \quad \forall(\mathbf{x}, t) \in \mathbb{R}^d \times \mathbb{R}^+, \end{aligned} \quad (5)$$

which are called normalization and nonnegativity constraints.

3 RELATED WORK

To better describe our algorithm and its connections with flow-based models, it is worthwhile to review prior research on normalizing flows.

3.1 NORMALIZING FLOWS AND CHANGE OF VARIABLES

Given a latent variable $\mathbf{z} \in \mathbb{Z} \subset \mathbb{R}^d$ drawn from a simple prior probability distribution p_Z . When normalizing flows transform \mathbf{z} into $\mathbf{x} = f(\mathbf{z}) \in \mathbb{X} \subset \mathbb{R}^d$ using a bijection $f: \mathbb{Z} \rightarrow \mathbb{X}$, the probability density function of \mathbf{x} follows the change of variables formula:

$$p_Z(\mathbf{z}) = p_X(\mathbf{x}) \left| \det \left(\frac{\partial f(\mathbf{z})}{\partial \mathbf{z}^T} \right) \right|, \quad (6)$$

where $\frac{\partial f(\mathbf{z})}{\partial \mathbf{z}^T}$ is the Jacobian of f at \mathbf{z} .

To describe a complex target distribution, a sequence of invertible and learnable mappings represented as $f_\theta = f_K \circ \dots \circ f_2 \circ f_1$ are constructed and θ denotes the trainable parameter of the neural network. Let $\mathbf{z}_1 = \mathbf{z}$, $\mathbf{z}_{k+1} = f_k(\mathbf{z}_k)$ ($k = 1, \dots, K$), then the log density of $\mathbf{x} = f_\theta(\mathbf{z})$ is computed by

$$\log p_X(\mathbf{x}) = \log p_Z(\mathbf{z}) - \sum_{k=1}^K \log \left| \det \left(\frac{\partial f_k(\mathbf{z}_k)}{\partial \mathbf{z}_k^T} \right) \right|. \quad (7)$$

There are some simple optional structures, such as planar flows and radial flows (Rezende & Mohamed, 2015). A key challenge is to increase the representative power of normalizing flows while simplifying the computation of the associated Jacobian determinant. NICE (Dinh et al., 2015) employed affine coupling layers, duplicating a portion of the input while transforming the remaining part, thereby maintaining reversibility. Real NVP (Dinh et al., 2016) introduced scale and translation parameters, further improving the performance of the flow. Additionally, the Jacobian of Real NVP transformations has a specific lower triangular structure, ensuring efficient computations with linear time complexity $\mathcal{O}(d)$ for the logdet-Jacobian term.

It is important to note that the resulting target distribution adheres to the probability normalization constraint, as guaranteed by the change of variables formula.

$$1 = \int_{\Omega_Z} p_Z(\mathbf{z}) d\mathbf{z} = \int p_X(\mathbf{x}) \left| \det \left(\frac{\partial f_\theta(\mathbf{z})}{\partial \mathbf{z}^T} \right) \right| d\mathbf{z} = \int_{\Omega_X} p_X(\mathbf{x}) d\mathbf{x} \quad (8)$$

3.2 CONTINUOUS NORMALIZING FLOWS AND INSTANTANEOUS CHANGE OF VARIABLES

Chen et al. (2018) proposed the Neural ODE framework and derived the finite normalizing flows to a continuous limit scheme, effectively expressing the invertible mappings using ODEs. The latent variable \mathbf{z} and its log probability change according to the instantaneous change of variables theorem (Villani, 2003):

$$\begin{cases} \frac{d\mathbf{z}(t)}{dt} = f_\theta(\mathbf{z}, t) \\ \frac{d \log p(\mathbf{z}(t), t)}{dt} = -\text{tr} \left(\frac{\partial f_\theta}{\partial \mathbf{z}} \right) = -\nabla \cdot f_\theta \end{cases} \quad (9)$$

By leveraging the adjoint sensitivity method to resolve augmented ODEs in reverse time, CNF computes gradients with respect to θ and is trained directly using maximum likelihood. An unexpected side-benefit is that the likelihood can be calculated using relatively cheap trace operations instead of the Jacobian determinant. Moreover, the entire transformation is automatically bijective when Eq.(9) has a unique solution. Therefore, the mapping f_θ does not need to be bijective.

3.3 RELATIONSHIP BETWEEN CONTINUOUS NORMALIZING FLOWS AND FOKKER-PLANCK EQUATIONS

Continuous normalizing flows can be related to the special case of the FP equation with zero diffusion, known as the Liouville equation. Let $\mathbf{z}(t) \in \mathbb{R}^d$ evolve through time according to the degenerate Eq.(1): $\frac{d\mathbf{z}(t)}{dt} = f(\mathbf{z}(t), t)$. As shown in Eq.(2), the probability density function $p(\mathbf{z}, t)$ satisfies the FP equation represented as:

$$\frac{\partial p(\mathbf{z}, t)}{\partial t} = -\nabla \cdot [p(\mathbf{z}, t)f(\mathbf{z}, t)]. \quad (10)$$

To compute the value of $p(\mathbf{z}, t)$, the characteristics method tracks the trajectory of a particle $\mathbf{z}(t)$. The total derivative of $p(\mathbf{z}(t), t)$ is given by

$$\frac{dp(\mathbf{z}(t), t)}{dt} = \frac{\partial p(\mathbf{z}, t)}{\partial t} + \frac{\partial p(\mathbf{z}, t)}{\partial \mathbf{z}} \cdot \frac{d\mathbf{z}}{dt} = -\nabla \cdot (pf) + \nabla p \cdot f = -p(\nabla \cdot f). \quad (11)$$

By dividing p on both sides, we obtain the same result as Eq.(9). (See Appendix A.2 in Chen et al. (2018) for more details).

4 PINF: PHYSICS-INFORMED NORMALIZING FLOWS

Let us begin by introducing the PINF algorithm for time-dependent FP equations, categorizing them into two scenarios based on whether the diffusion term is zero. Subsequently, we will present the special design of the algorithm for solving the steady-state FP equation.

4.1 TIME-DEPENDENT FOKKER-PLANCK EQUATIONS

The TFP equation is essentially an initial value problem (2), where the initial density function is known.

4.1.1 TFP EQUATION WITH ZERO DIFFUSION

Problem Setup:

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla \cdot [p(\mathbf{x}, t)\mu(\mathbf{x}, t)] \\ p(\mathbf{x}, 0) &= p_0(\mathbf{x}) \end{aligned} \quad (12)$$

The objective of solving the TFP equation is to compute the density p at any given point (\mathbf{x}', t') . Following the derivation in Eq.(11), the TFP equation can be reformulated as an initial value problem of ODEs.

$$\begin{cases} \frac{d\mathbf{x}(t)}{dt} = \mu(\mathbf{x}, t), & \mathbf{x}(t') = \mathbf{x}' \\ \frac{d \log p(\mathbf{x}(t), t)}{dt} = -\nabla \cdot \mu(\mathbf{x}, t) \end{cases} \quad (13)$$

Notably, the increment in $\log p$ is independent of its value and solely depends on the start time t' , the stop time $t_0 = 0$, and the value of \mathbf{x} . Therefore, the initial states for $(\mathbf{x}, \log p)$ can be set as $(\mathbf{x}', 0)$, and the solution is obtained by using the corresponding output of ODE solvers, yielding $(\mathbf{x}_0, \Delta \log p)$. The specific algorithm is outlined as follows:

Algorithm 1 PINF algorithm for TFP equations with zero diffusion

Input: drift term $\mu(\mathbf{x}, t)$, samples \mathbf{x}' , time t' , initial PDF $p_0(\mathbf{x})$.
def $f_{aug}([\mathbf{x}_t, \log p_t], t)$: ▷ ODEs dynamics
 return $[\mu, -\nabla \cdot \mu]$ ▷ Concatenate dynamics of state and log-density
 $[\mathbf{x}_0, \Delta \log p] \leftarrow \text{ODESolve}(f_{aug}, [\mathbf{x}', 0], t', 0)$ ▷ Calculate $\int_{t'}^0 f_{aug}([\mathbf{x}(t), \log p(\mathbf{x}(t), t)], t) dt$
 $\log \hat{p} \leftarrow \log p_0(\mathbf{x}_0) - \Delta \log p$ ▷ Add change in log-density
 $\hat{p}(\mathbf{x}', t') = e^{\log \hat{p}}$
Output: $\hat{p}(\mathbf{x}', t')$

The key distinction between PINF and CNF lies in the fact that the drift $\mu(\mathbf{x}, t)$ is known in the TFP equation. In this context, there is no need to train f_θ from the real data samples and we can simply use ODE solvers once to obtain the outcomes.

4.1.2 TFP EQUATION WITH DIFFUSION

Problem Setup:

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla \cdot [p(\mathbf{x}, t)\mu(\mathbf{x}, t)] + \nabla \cdot [\nabla \cdot (p(\mathbf{x}, t)\mathbf{D}(\mathbf{x}, t))] \\ p(\mathbf{x}, 0) &= p_0(\mathbf{x}) \end{aligned} \quad (14)$$

We first perform some necessary transformations of the equation.

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla \cdot (p\mu) + \nabla \cdot [\nabla \cdot (p\mathbf{D})] \\ &= -\nabla \cdot [p\mu - \nabla \cdot (p\mathbf{D})] \\ &= -\nabla \cdot [p\mu - (\nabla p)\mathbf{D} - p(\nabla \cdot \mathbf{D})] \\ &= -\nabla \cdot [p(\mu - (\nabla \log p)\mathbf{D} - \nabla \cdot \mathbf{D})] \\ &= -\nabla \cdot (p\mu^*) \end{aligned} \quad (15)$$

Let us define $\mu^*(\mathbf{x}, t) := \mu - (\nabla \log p)\mathbf{D} - \nabla \cdot \mathbf{D}$. We force \mathbf{x} to evolve through time following $\frac{d\mathbf{x}}{dt} = \mu^*(\mathbf{x}, t)$ and get the associated ODEs by Eq.(11).

$$\begin{cases} \frac{d\mathbf{x}(t)}{dt} = \mu^*(\mathbf{x}, t) \\ \frac{d \log p(\mathbf{x}(t), t)}{dt} = -\nabla \cdot \mu^*(\mathbf{x}, t) \end{cases} \quad (16)$$

In particular, unlike the zero diffusion case, the ODE dynamics depend on the unknown function $\log p$ and its gradient will not be evaluated during the forward calculation. For this reason, we parameterize $\log p(\mathbf{x}, t)$ as a neural network ϕ_θ .

Neural network architecture. The network structure is inspired by the value function representations used for stochastic optimal control (Li et al., 2022), multi-agent optimal control (Onken et al., 2021b), and high-dimensional optimal control (Onken et al., 2022). The network is given by

$$u(\mathbf{s}; \theta) = \mathbf{w}^\top \mathcal{N}(\mathbf{s}; \theta_{\mathcal{N}}) + \frac{1}{2} \mathbf{s}^\top (\mathbf{A}^\top \mathbf{A}) \mathbf{s} + \mathbf{b}^\top \mathbf{s} + c, \quad (17)$$

where $\theta = (\mathbf{w}, \theta_{\mathcal{N}}, \mathbf{A}, \mathbf{b}, c)$ are the trainable weights of u_θ . The shapes of those parameters are listed as follows: the inputs $\mathbf{s} = (\mathbf{x}, t) \in \mathbb{R}^{d+1}$ corresponding to space-time, $\mathcal{N}(\mathbf{s}; \theta_{\mathcal{N}}) : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^m$, $\mathbf{w} \in \mathbb{R}^m$, $\mathbf{A} \in \mathbb{R}^{r \times (d+1)}$, $\mathbf{b} \in \mathbb{R}^{d+1}$, and $c \in \mathbb{R}$. The rank $r = \min(10, d+1)$ is set to limit the number of parameters in $\mathbf{A}^\top \mathbf{A}$. Here, \mathbf{A} , \mathbf{b} , and c model quadratic potentials, i.e., linear dynamics; \mathcal{N} models nonlinear dynamics.




Figure 1: An illustration of the PINF algorithm for TFP equations (left) and SFP equations (right).

In our implementation, for \mathcal{N} , we use a residual neural network (ResNet) (He et al., 2016) with $(L + 1)$ layers and obtain $\mathcal{N}(s; \theta_{\mathcal{N}}) = \mathbf{a}_L$ as the final step of the forward propagation

$$\begin{aligned} \mathbf{a}_0 &= \sigma(\mathbf{K}_0 \mathbf{s} + \mathbf{b}_0), \\ \mathbf{a}_k &= \mathbf{a}_{k-1} + h \sigma(\mathbf{K}_k \mathbf{a}_{k-1} + \mathbf{b}_k), \quad (1 \leq k \leq L) \end{aligned} \quad (18)$$

where the trainable weights are $\mathbf{K}_0 \in \mathbb{R}^{m \times (d+1)}$, $\mathbf{K}_k \in \mathbb{R}^{m \times m}$ ($1 \leq k \leq L$), $\mathbf{b}_k \in \mathbb{R}^m$ ($0 \leq k \leq L$), and $\theta_{\mathcal{N}} = \{(K_k, b_k)\}_{k=0}^L$. We select the step size $h = 1/L$ and the element-wise activation function $\sigma(x) = \log(\exp(x) + \exp(-x))$ (Onken et al., 2021a), which is the antiderivative of the hyperbolic tangent, i.e., $\sigma'(x) = \tanh(x)$. It can also be seen as a smoothed absolute value function (Ruthotto et al., 2020).

To satisfy the initial value condition $p(\mathbf{x}, 0) = p_0(\mathbf{x})$, the network ϕ_θ is cast as

$$\phi_\theta(\mathbf{x}, t) = \log p_0(\mathbf{x}) + tu(\mathbf{x}, t; \theta), \quad (19)$$

to represent $\log p(\mathbf{x}, t)$ (Lagaris et al., 1998). This structured design can impose hard constraints that are strictly enforced, rather than soft constraints (Wang & Yu, 2021).

Self-supervised training method. Distinct from normalizing flows that minimize the Kullback-Leibler divergence, we employ a self-supervised training method. Our approach eliminates the need for labeled data or real data samples from the target distribution, which are sometimes challenging or costly to acquire. Training data is adaptively generated based on the initial PDF.

For training ϕ_θ , we solve the ODEs (16) related to the network ϕ_θ , obtaining predictions in the form of $\log p_{ode}$. Moreover, the network can also output the prediction $\log p_{net}$ directly. We calculate the Mean Squared Error (MSE) between $\log p_{ode}$ and $\log p_{net}$ and use the Adam optimizer (Kingma & Ba, 2014).

$$\mathcal{L} = \frac{1}{N} \sum_{k=1}^N \|\log p_{ode}(\mathbf{x}_k, t_k) - \log p_{net}(\mathbf{x}_k, t_k)\|^2 \quad (20)$$

To address issues such as small gradients and optimization difficulties in high-dimensional scenarios, we avoid comparing p_{pred} directly in the loss function: $\mathcal{L} = \text{MSE}(p_{ode}, p_{net})$. Experimental results also demonstrate that applying the logarithm function accelerates the training process.

Flexible prediction modes. With the trained ϕ_θ , we provide two flexible prediction modes: numerical solutions solved by ODE solvers and continuous solutions predicted by neural networks. This flexibility enables a trade-off between efficiency and accuracy, maintaining advantages such as memory efficiency, adaptive solvers, and parallel computation (Kang et al., 2021). When using neural networks for prediction, our approach is mesh-free and causality-free (Nakamura-Zimmerer et al., 2020), similar to PINNs. Data at different time steps can be computed rapidly and in parallel, allowing for real-time or low-power applications. Alternatively, when using ODE solvers for prediction, we can freely choose from modern ODE solvers for adaptive computation (Runge, 1895; Wanner & Hairer, 1996). See Alg.(2) for the complete algorithm.

Algorithm 2 PINF algorithm for TFP equations with diffusion

Input: drift term $\mu(\mathbf{x}, t)$, diffusion matrix $\mathbf{D}(\mathbf{x}, t)$, samples \mathbf{x} , time t , initial PDF $p_0(\mathbf{x})$, stop time T , maximum iteration number M , learning rate η .

def $f_{aug}([\mathbf{x}_t, \log p_t], t)$: ▷ ODEs dynamics
 $\mu^* \leftarrow \mu - (\nabla \phi_\theta) \mathbf{D} - \nabla \cdot \mathbf{D}$ ▷ Characteristic curves $\frac{d\mathbf{x}(t)}{dt}$
return $[\mu^*, -\nabla \cdot \mu^*]$ ▷ Concatenate dynamics of state and log-density

Train ϕ_θ :
for $k = 0, \dots, M$:
Uniformly sample $t_k \in [0, T]$ ▷ Sample training data
Sample mini-batch $\mathbf{x}_0^k \sim p_0(\mathbf{x})$
 $[\mathbf{x}^k, \log p(\mathbf{x}^k, t_k)] \leftarrow \text{ODESolve}(f_{aug}, [\mathbf{x}_0^k, \log p_0(\mathbf{x}_0^k)], 0, t_k)$
▷ Calculate $[\mathbf{x}_0^k, \log p_0(\mathbf{x}_0^k)] + \int_0^{t_k} f_{aug}([\mathbf{x}(t), \log p(\mathbf{x}(t), t)], t) dt$
ODEs prediction: $\log p_{ode} \leftarrow \log p(\mathbf{x}^k, t_k)$ ▷ Two prediction modes
 ϕ_θ prediction: $\log p_{net} \leftarrow \phi_\theta(\mathbf{x}^k, t_k)$
Compute the MSE loss: $\mathcal{L} = \text{MSE}(\log p_{ode}, \log p_{net})$ ▷ Calculate loss on mini-batch \mathbf{x}^k
Update the parameters θ using the Adam optimizer with learning rate η . ▷ Train ϕ_θ once

Predict \hat{p} :
 ϕ_θ mode: $\hat{p}_{net}(\mathbf{x}, t) = e^{\phi_\theta(\mathbf{x}, t)}$
ODEs mode: $[\mathbf{x}_0, \Delta \log p] \leftarrow \text{ODESolve}(f_{aug}, [\mathbf{x}, 0], t, 0)$
 $\log \hat{p}_{ode} \leftarrow \log p_0(\mathbf{x}_0) - \Delta \log p$ ▷ Add change in log-density
 $\hat{p}_{ode}(\mathbf{x}, t) = e^{\log \hat{p}_{ode}}$

Output: $\hat{p}_{net}(\mathbf{x}, t), \hat{p}_{ode}(\mathbf{x}, t)$

4.2 STEADY-STATE FOKKER-PLANCK EQUATIONS

Problem Setup:

$$\begin{aligned} \frac{\partial p(\mathbf{x}, t)}{\partial t} &= -\nabla \cdot [p(\mathbf{x}, t)\mu(\mathbf{x}, t)] + \nabla \cdot [\nabla \cdot (p(\mathbf{x}, t)\mathbf{D}(\mathbf{x}, t))] \\ \frac{\partial p(\mathbf{x}, t)}{\partial t} &= 0 \end{aligned} \quad (21)$$

The SFP equation replaces the initial value condition with an equation constraint (3), ensuring that its solution remains invariant over time. Therefore, the steady-state PDF only depends on spatial variables \mathbf{x} and we derive another form of our problem:

$$-\nabla \cdot [p(\mathbf{x})\mu] + \nabla \cdot [\nabla \cdot (p(\mathbf{x})\mathbf{D})] = 0 \quad (22)$$

Normalization Challenge. In the TFP equations, since the initial PDF $p_0(\mathbf{x})$ satisfies the normalization constraint, the total density integral on the space domain $\int_{\mathbb{R}^d} p(\mathbf{x}, t) d\mathbf{x}$ remains conserved as the solution evolves under equation constraints or ODEs controls. Consequently, the solution $p(\mathbf{x}, t)$ predicted by the PINF algorithm also satisfies the normalization constraint.

However, for the SFP equations, it is easily checked that if $p(\mathbf{x})$ is the solution, then any $cp(\mathbf{x})$ ($c > 0$ is a constant) also satisfies the equation. In the ODEs, the scheme remains unchanged due to $\nabla \log(cp(\mathbf{x})) = \nabla(\log p(\mathbf{x}) + \log c) = \nabla \log p(\mathbf{x})$. Hence, starting from \mathbf{x}_1 and evolving through t_1 to t_2 , both ODEs associated with $p(\mathbf{x})$ and $cp(\mathbf{x})$ reach the same point \mathbf{x}_2 and have the same increment $\Delta \log p$. For the ODEs of $cp(\mathbf{x})$, the predicted value of $\log \hat{p}(\mathbf{x}_2)$ is given by

$$\begin{aligned} \log \hat{p}(\mathbf{x}_2) &= \log(cp(\mathbf{x}_1)) + \Delta \log p \\ &= \log c + (\log p(\mathbf{x}_1) + \Delta \log p) \\ &= \log c + \log p(\mathbf{x}_2) \\ &= \log(cp(\mathbf{x}_2)) \end{aligned} \quad (23)$$

Ultimately, $p(\mathbf{x})$ and $cp(\mathbf{x})$ satisfy the ODEs.

Neural network architecture. We still employ a neural network $\phi_\theta(\mathbf{x})$ to parameterize $\log p(\mathbf{x})$ in ODEs (16). For the normalization constraint, we adopt the Real NVP (Dinh et al., 2016) to model the equilibrium distribution from a Gaussian distribution $p_G(\mathbf{z}; \mathbf{0}, \mathbf{I})$. We stack L affine coupling

layers to build a flexible and tractable bijective transformation. At each layer, given the input $\mathbf{x} \in \mathbb{R}^d$ and dimension $n < d$, the output \mathbf{y} is defined as

$$\begin{aligned} \mathbf{y}_{1:n} &= \mathbf{x}_{1:n} \\ \mathbf{y}_{n+1:d} &= \mathbf{x}_{n+1:d} \odot \exp(s(\mathbf{x}_{1:n})) + t(\mathbf{x}_{1:n}) \end{aligned} \quad (24)$$

where s and t represent scale and translation, and \odot is the element-wise product. The Jacobian of this transformation is triangular, reads

$$\frac{\partial \mathbf{y}}{\partial \mathbf{x}^T} = \begin{bmatrix} \mathbf{1}_{n \times n} & 0 \\ \frac{\partial \mathbf{y}_{n+1:d}}{\partial \mathbf{x}_{1:n}^T} & \text{diag}(\exp[s(\mathbf{x}_{1:n})]) \end{bmatrix} \quad (25)$$

Consequently, we can efficiently compute its log-determinant as $\sum_k s(\mathbf{x}_{1:n})_k$.

Since partitioning can be implemented using a binary mask b , the forward and backward computation processes respectively follow the equations,

$$\begin{aligned} \text{Forward: } \mathbf{x}_b &= b \odot \mathbf{x} \\ \mathbf{y} &= \mathbf{x}_b + (1 - b) \odot (\mathbf{x} \odot \exp(s(\mathbf{x}_b)) + t(\mathbf{x}_b)) \\ \log \left| \det \left(\frac{\partial \mathbf{y}}{\partial \mathbf{x}^T} \right) \right| &= \sum (1 - b) \odot s(\mathbf{x}_b) \\ \text{Backward: } \mathbf{y}_b &= b \odot \mathbf{y} \\ \mathbf{x} &= \mathbf{y}_b + (1 - b) \odot (\mathbf{y} - t(\mathbf{y}_b)) \odot \exp(-s(\mathbf{y}_b)) \\ \log \left| \det \left(\frac{\partial \mathbf{x}}{\partial \mathbf{y}^T} \right) \right| &= - \sum (1 - b) \odot s(\mathbf{y}_b) \end{aligned} \quad (26)$$

Here, $s, t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ are constructed via a simple 3-layer MLP (Goodfellow et al., 2016) with the activation function $\sigma(x) = \tanh(x)$.

$$\begin{aligned} \mathbf{a}_1 &= \sigma(\mathbf{K}_1 \mathbf{x} + \mathbf{b}_1), \\ \mathbf{a}_2 &= \sigma(\mathbf{K}_2 \mathbf{a}_1 + \mathbf{b}_2), \\ \mathbf{a}_3 &= \mathbf{K}_3 \mathbf{a}_2 + \mathbf{b}_3. \end{aligned} \quad (27)$$

The self-supervised training method and loss function remain consistent with the previous description. This is the PINF algorithm for SFP equations.

Algorithm 3 PINF algorithm for SFP equations

Input: drift term μ , diffusion matrix \mathbf{D} , samples \mathbf{x} , maximum iteration number M , learning rate η .

def $f_{aug}([\mathbf{x}_t, \log p_t], t)$: ▷ ODEs dynamics
 $\mu^* \leftarrow \mu - (\nabla \phi_\theta) \mathbf{D} - \nabla \cdot \mathbf{D}$ ▷ Characteristic curves $\frac{d\mathbf{x}(t)}{dt}$
return $[\mu^*, -\nabla \cdot \mu^*]$ ▷ Concatenate dynamics of state and log-density

Train ϕ_θ :

for $k = 0, \dots, M$:

Sample mini-batch $\mathbf{z}_0^k \sim p_G(\mathbf{z})$ ▷ Sample training data

Forward computation: $\mathbf{x}_0^k, \log p(\mathbf{x}_0^k) = \phi_\theta^f(\mathbf{z}_0^k, \log p_G(\mathbf{z}_0^k))$

$[\mathbf{x}_1^k, \log p(\mathbf{x}_1^k)] \leftarrow \text{ODESolve}(f_{aug}, [\mathbf{x}_0^k, \log p(\mathbf{x}_0^k)], 0, 1)$

▷ Calculate $[\mathbf{x}_0^k, \log p(\mathbf{x}_0^k)] + \int_0^1 f_{aug}([\mathbf{x}(t), \log p(\mathbf{x}(t), t)], t) dt$

ODEs prediction: $\log p_{\text{ode}} \leftarrow \log p(\mathbf{x}_1^k)$

ϕ_θ prediction: $\log p_{\text{net}} \leftarrow \phi_\theta^b(\mathbf{x}_1^k)$

Compute the MSE loss: $\mathcal{L} = \text{MSE}(\log p_{\text{ode}}, \log p_{\text{net}})$ ▷ Calculate loss on mini-batch \mathbf{x}_1^k

Update the parameters θ using the Adam optimizer with learning rate η . ▷ Train ϕ_θ once

Predict \hat{p} :

$\hat{p}_{\text{net}}(\mathbf{x}) = e^{\phi_\theta^b(\mathbf{x})}$

Output: $\hat{p}_{\text{net}}(\mathbf{x})$

5 EXPERIMENT

In this section, we consider the following three numerical examples to test the performance of our PINF algorithm.

5.1 A TOY EXAMPLE

We first present a toy example, a TFP equation with zero diffusion. We employ it to illustrate that our PINF algorithm is the method of characteristics in mathematics.

$$\begin{aligned} \frac{\partial p}{\partial t} + 2t\nabla \cdot (p\mathbf{1}) &= 0, \quad t \in \mathbb{R}^+, \mathbf{x} \in \mathbb{R}^d \\ p(\mathbf{x}, 0) &= (2\pi)^{-d/2} \exp\left(-\frac{1}{2}\|\mathbf{x} + \mathbf{1}\|^2\right) \end{aligned} \quad (28)$$

Here $\mathbf{1} \in \mathbb{R}^d$ denotes an all-ones vector, and the diffusion term $\mu = 2t \cdot \mathbf{1}$ is independent of \mathbf{x} . The corresponding ODEs obtained through the PINF algorithm (1) are as follows:

$$\begin{cases} \frac{d\mathbf{x}(t)}{dt} = 2t \cdot \mathbf{1} \\ \frac{d \log p(\mathbf{x}(t), t)}{dt} = 0 \end{cases} \quad (29)$$

Consequently, the analytical solution of the equation is

$$p(\mathbf{x}, t) = p(\mathbf{x}_0, 0) = p_0(\mathbf{x} - t^2 \cdot \mathbf{1}) = (2\pi)^{-d/2} \exp\left(-\frac{1}{2}\|\mathbf{x} + (1 - t^2)\mathbf{1}\|^2\right) \quad (30)$$

5.2 TFP EQUATION

We consider a relatively high-dimensional TFP equation. For this problem, the PINN is effective primarily with the dimension $d \leq 3$.

$$\begin{aligned} \frac{\partial p}{\partial t} - \frac{1}{2}\Delta p + 2\nabla \cdot (p\mathbf{1}) &= 0, \quad t \in [0, 1], \mathbf{x} \in \mathbb{R}^d \\ p(\mathbf{x}, 0) &= (2\pi)^{-d/2} \exp(\|\mathbf{x}\|^2/2) \end{aligned} \quad (31)$$

The exact solution is given by

$$p(\mathbf{x}, t) = \frac{1}{(2\pi(t+1))^{d/2}} \exp\left(-\frac{\|\mathbf{x} - 2t \cdot \mathbf{1}\|^2}{2(t+1)}\right) \quad (32)$$

Using the PINF algorithm (2), we formulate the ODEs as

$$\begin{cases} \frac{d\mathbf{x}(t)}{dt} = 2 \cdot \mathbf{1} - \frac{1}{2}\nabla \log p \\ \frac{d \log p(\mathbf{x}(t), t)}{dt} = -\nabla \cdot \left(\frac{d\mathbf{x}}{dt}\right) \end{cases} \quad (33)$$

We solve this TFP equation of $d = 10$. We take $L = 4$ residual layers with $m = 32$ hidden neurons and set training iteration $M = 10000$, the learning rate for Adam optimizer $\eta = 0.01$, and the batch size is 2000. The initial spatial training set is generated from $p_0(\mathbf{x})$, and the corresponding temporal training set is uniformly sampled in the interval $[0, T]$. We evaluate the performance of our PINF algorithm at the point $(\mathbf{x}_1, \mathbf{x}_2, 2, \dots, 2) \in \mathbb{R}^d$ and the stop time $T = 1$, where $(\mathbf{x}_1, \mathbf{x}_2)$ is drawn from a uniform grid within the finite spatial domain $[-5, 5]^2$. The comparison between the prediction and the ground truth is shown in Fig.(2) and we observe a good agreement.




Figure 2: Exact and predicted solutions for the TFP equation from Case 2 at the stop time $T = 1$. Right panel: MAPE between the exact solution and the prediction p_{ode} .

5.3 HIGH DIMENSIONAL SFP EQUATION

In this part, a SFP equation with drift term $\mu(\mathbf{x}) = -a\mathbf{x}$ and diffusion matrix $\mathbf{D} = \frac{\sigma^2}{2}\mathbf{1}_{d \times d}$ is considered.

$$-\nabla \cdot (p(\mathbf{x})\mu) + \nabla \cdot [\nabla \cdot (p(\mathbf{x})\mathbf{D})] = 0, \quad \mathbf{x} \in \mathbb{R}^d \tag{34}$$

Its exact solution is a single Gaussian distribution, represented as

$$p(\mathbf{x}) = \left(\frac{a}{\pi\sigma^2}\right)^{d/2} \exp\left(-\frac{a\|\mathbf{x}\|^2}{\sigma^2}\right) \tag{35}$$

In our experiment, we set $a = \sigma = 1$ and evaluate our PINF algorithm on the SFP equation with dimensions $d = 30$ and $d = 50$. We employ $L = 4$ affine coupling layers and set $M = 500$, $\eta = 0.01$, batch size is 2000. Since its steady-state solution is an unimodal function near the origin, we select the test points as $(\mathbf{x}_1, \mathbf{x}_2, 0, \dots, 0) \in \mathbb{R}^d$, where $(\mathbf{x}_1, \mathbf{x}_2)$ is drawn from a uniform grid within the range $[-3, 3]^2$, resulting in total $50 \times 50 = 2,500$ test points. Fig.(3) shows the exact solution $p(\mathbf{x})$ and our PINF solution $p(\mathbf{x}; \theta) = \phi_\theta^b(\mathbf{x})$, where it can be seen that they are visually indistinguishable, and the relative error remains below 0.2%, even for relatively high dimensions.

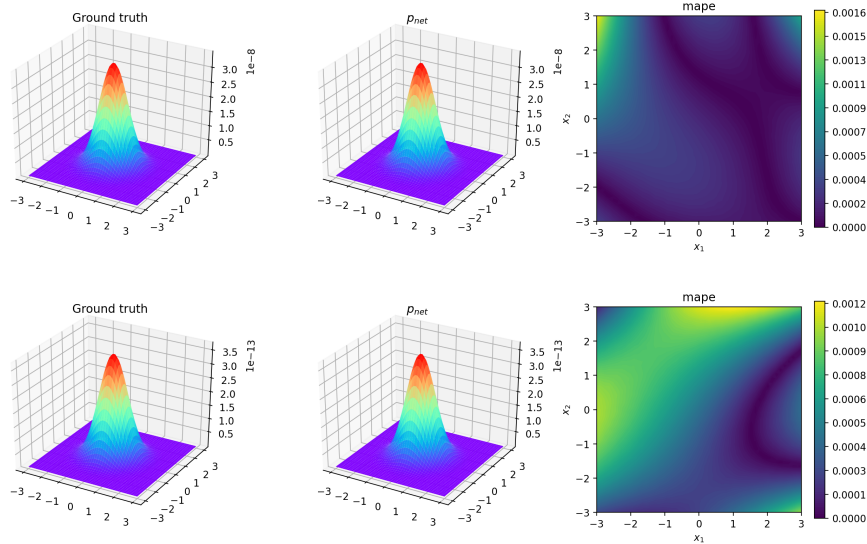


Figure 3: Exact and predicted solutions for the SFP equation from Case 3. Top row: $d=30$. Bottom row: $d=50$.

6 DISCUSSION

Our proposed PINF algorithm extends CNF to the scheme with diffusion using the method of characteristics. Compared to PINNs, which apply pointwise equation constraints, our improvements are mainly in reformulating the FP equation as ODEs and computing the integral of the loss along characteristic curves using an ODE solver, thereby enhancing training efficiency and stability. It effectively addresses the normalization constraint of the PDF even in high-dimensional FP problems.

The critical distinction between PINF for FP equations and CNF for density estimation lies in whether the drift term μ is known. In FP equations, the drift term $\mu(\mathbf{x}, t)$ is known and we obtain its solution by self-supervised training of a neural network. Conversely, the drift term μ_θ necessitates likelihood-based training for the latter. In future work, we may explore the application of PINF with diffusion for density estimation tasks to enhance model generalization and data generation quality.

Considering the profound physical significance and wide-ranging applications of FP equations, the PINF algorithm can also be seamlessly integrated with large-scale mean-field games for optimal control problems, such as path planning, obstacle avoidance, and tracking in drone swarm applications.

REFERENCES

- Michael S Albergo, Gurtej Kanwar, and Phiala E Shanahan. Flow-based generative models for markov chain monte carlo in lattice field theory. *Physical Review D*, 100(3):034515, 2019.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Weihua Deng. Finite element method for the space and time fractional fokker-planck equation. *SIAM journal on numerical analysis*, 47(1):204–226, 2009.
- L Dinh, D Krueger, and Y Bengio. Nice: non-linear independent components estimation. In *International Conference on Learning Representations Workshop*, 2015.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. In *International Conference on Learning Representations*, 2016.
- Xiaodong Feng, Li Zeng, and Tao Zhou. Solving time dependent fokker-planck equations via temporal normalizing flow. *Communications in Computational Physics*, 32(2):401–423, 2022.
- Chris Finlay, Jörn-Henrik Jacobsen, Levon Nurbekyan, and Adam Oberman. How to train your neural ode: the world of jacobian and kinetic regularization. In *International conference on machine learning*, pp. 3154–3164. PMLR, 2020.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pp. 2722–2730. PMLR, 2019.
- Wei Kang, Qi Gong, Tenavi Nakamura-Zimmerer, and Fariba Fahroo. Algorithms of data generation for deep learning and feedback design: A survey. *Physica D: Nonlinear Phenomena*, 425:132955, 2021.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.
- Pankaj Kumar and S Narayanan. Solution of fokker-planck equation by finite element and finite difference methods for nonlinear systems. *Sadhana*, 31:445–461, 2006.
- Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.
- Xingjian Li, Deepanshu Verma, and Lars Ruthotto. A neural network approach for stochastic optimal control. *arXiv preprint arXiv:2209.13104*, 2022.
- Tenavi Nakamura-Zimmerer, Qi Gong, and Wei Kang. A causality-free neural network method for high-dimensional hamilton-jacobi-bellman equations. In *2020 American Control Conference (ACC)*, pp. 787–793. IEEE, 2020.
- Bernt Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- Derek Onken, Samy Wu Fung, Xingjian Li, and Lars Ruthotto. Ot-flow: Fast and accurate continuous normalizing flows via optimal transport. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(10):9223–9232, 2021a.

- Derek Onken, Levon Nurbekyan, Xingjian Li, Samy Wu Fung, Stanley Osher, and Lars Ruthotto. A neural network approach applied to multi-agent optimal control. In *2021 European Control Conference (ECC)*, pp. 1036–1041. IEEE, 2021b.
- Derek Onken, Levon Nurbekyan, Xingjian Li, Samy Wu Fung, Stanley Osher, and Lars Ruthotto. A neural network approach for high-dimensional optimal control applied to multi-agent path finding. *IEEE Transactions on Control Systems Technology*, 31(1):235–251, 2022.
- 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.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pp. 1530–1538. PMLR, 2015.
- Hannes Risken. *Fokker-planck equation*. Springer, 1996.
- Carl Runge. Über die numerische auflösung von differentialgleichungen. *Mathematische Annalen*, 46(2):167–178, 1895.
- Lars Ruthotto, Stanley J Osher, Wuchen Li, Levon Nurbekyan, and Samy Wu Fung. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183–9193, 2020.
- Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375:1339–1364, 2018.
- Kejun Tang, Xiaoliang Wan, and Qifeng Liao. Adaptive deep density approximation for fokker-planck equations. *Journal of Computational Physics*, 457:111080, 2022.
- Cédric Villani. *Topics in Optimal Transportation*. Number 58. American Mathematical Soc., 2003.
- Rui Wang and Rose Yu. Physics-guided deep learning for dynamical systems: A survey. *arXiv preprint arXiv:2107.01272*, 2021.
- Gerhard Wanner and Ernst Hairer. *Solving ordinary differential equations II*, volume 375. Springer Berlin Heidelberg New York, 1996.
- M Fo Wehner and WG Wolfer. Numerical evaluation of path-integral solutions to fokker-planck equations. *Physical Review A*, 27(5):2663, 1983.
- E Weinan and Bing Yu. The deep ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6(1):1–12, 2018.
- Liu Yang and George Em Karniadakis. Potential flow generator with l_2 optimal transport regularity for generative models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2): 528–538, 2020.
- Linfeng Zhang, Lei Wang, et al. Monge-ampère flow for generative modeling. *arXiv e-prints*, pp. arXiv–1809, 2018.