



Data Seminar

04/02/2021



Outline

1. DeepXDE: A Deep Learning Library of Solving Differential Equations. Lu L, Meng X, Mao Z, et al. SIAM Review 2021.
2. Solving high-dimensional partial differential equations using deep learning. Han J, Jentzen A, Weinan E. PNAS 2018.
3. Forecasting Sequential Data Using Consistent Koopman Autoencoders. Azencot O, Erichson N B, Lin V, et al. ICML 2020.



Outline

1. **DeepXDE: A Deep Learning Library of Solving Differential Equations. Lu L, Meng X, Mao Z, et al. SIAM Review 2021.**
2. Solving high-dimensional partial differential equations using deep learning. Han J, Jentzen A, Weinan E. PNAS 2018.
3. Forecasting Sequential Data Using Consistent Koopman Autoencoders. Azencot O, Erichson N B, Lin V, et al. ICML 2020.



Motivation & Idea

Motivation:

- Traditional mesh-based methods to solve PDEs like FDM and FEM: Suffer from the curse of dimensionality.
- Deep learning method: Mesh free, taking advantage of automatic differentiation.

Ideas:

- Physics-Informed Neural Network (PINN) to embed a PDE into the loss of the neural network.
- Residual-based adaptive refinement (RAR) to improve the training efficiency.

Loss function

The loss is composed by two parts: PDE equation loss and boundary loss:

- $\mathcal{L}_f(\theta, \mathcal{T}_f)$ is the PDE equation loss
 - \mathcal{T}_f are the points \mathbf{x} that we will calculate the $\|f(\mathbf{x}, \lambda)\|^2$, but $u(\mathbf{x}, t)$ for \mathbf{x} in \mathcal{T}_f is not known.
- $\mathcal{L}_b(\theta, \mathcal{T}_b)$ is the boundary loss
 - \mathcal{T}_b are the points \mathbf{x} that we will calculate the boundary loss, where we know $B(u, \mathbf{x})=0$.
- w_f and w_b are loss weights.

$$(2.2) \quad \mathcal{L}(\theta; \mathcal{T}) = w_f \mathcal{L}_f(\theta; \mathcal{T}_f) + w_b \mathcal{L}_b(\theta; \mathcal{T}_b),$$

where

$$\mathcal{L}_f(\theta; \mathcal{T}_f) = \frac{1}{|\mathcal{T}_f|} \sum_{\mathbf{x} \in \mathcal{T}_f} \left\| f \left(\mathbf{x}; \frac{\partial \hat{u}}{\partial x_1}, \dots, \frac{\partial \hat{u}}{\partial x_d}; \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_1}, \dots, \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_d}; \dots; \boldsymbol{\lambda} \right) \right\|_2^2,$$

$$\mathcal{L}_b(\theta; \mathcal{T}_b) = \frac{1}{|\mathcal{T}_b|} \sum_{\mathbf{x} \in \mathcal{T}_b} \|B(\hat{u}, \mathbf{x})\|_2^2,$$

Solving Inverse Problems: Infer λ

This time, some PDE parameters λ are unknown, but we know some $I(u, \mathbf{x}) = 0$ for some points \mathbf{x} .

Extra Loss (data loss):

We add another loss term $\mathcal{L}_i(\theta, \lambda, \mathcal{T}_i)$ where \mathcal{T}_i are the points we will calculate data loss, where we know some extra information that $I(u, \mathbf{x}) = 0$.

$$\mathcal{L}(\theta, \lambda; \mathcal{T}) = w_f \mathcal{L}_f(\theta, \lambda; \mathcal{T}_f) + w_b \mathcal{L}_b(\theta, \lambda; \mathcal{T}_b) + w_i \mathcal{L}_i(\theta, \lambda; \mathcal{T}_i),$$

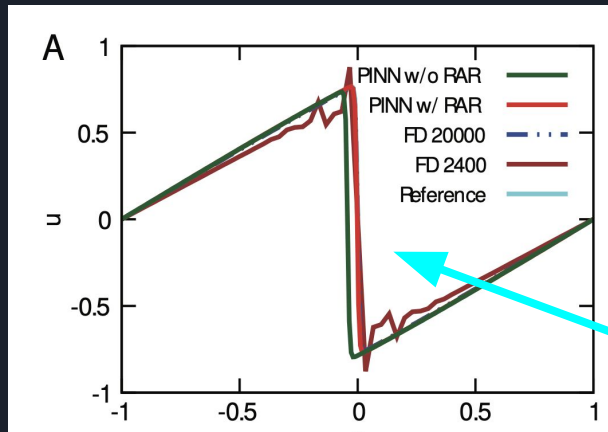
where

$$\mathcal{L}_i(\theta, \lambda; \mathcal{T}_i) = \frac{1}{|\mathcal{T}_i|} \sum_{\mathbf{x} \in \mathcal{T}_i} \|\mathcal{I}(\hat{u}, \mathbf{x})\|_2^2.$$

Residual-Based Adaptive Refinement (RAR)

However, some PDE has steep gradients (discontinuity).

We add more residual points in the locations where PDE equation loss is large until the mean PDE loss is smaller than a pre-set threshold ϵ_0 .



$$(2.3) \quad \mathcal{E}_r = \frac{1}{V} \int_{\Omega} \left\| f \left(\mathbf{x}; \frac{\partial \hat{u}}{\partial x_1}, \dots, \frac{\partial \hat{u}}{\partial x_d}; \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_1}, \dots, \frac{\partial^2 \hat{u}}{\partial x_1 \partial x_d}; \dots; \boldsymbol{\lambda} \right) \right\| dx$$

is smaller than a threshold \mathcal{E}_0 , where V is the volume of Ω .

Steep gradients (sudden change of u)

V : space of x :
Eg. $0 < x_1 < 1, 0 < x_2 < 3$, then $V=3$

Experiment 1: PINN vs SEM

In experiment 1, the author wants to show that PINN solved $u(x,t)$ is close to the SEM solution (Here SEM solution is seen as ground-truth.)





Fig. 7 Example in subsection 4.1. Comparison of the PINN solution with the solution obtained by using the spectral element method (SEM). (A) The SEM solution u_{SEM} , (B) the PINN solution u_{NN} , (C) the absolute error $|u_{SEM} - u_{NN}|$.

Experiment 1: PINN vs SEM




Experiment 2: PINN With RAR

In experiment 2, the author wants to show that PINN w/ RAR fits the ground-truth better than PINN w/o RAR.

- Reference is ground truth.
- FD (Finite difference method) is the baseline.
- FD 20000/2400 means 20000/2400 mesh.

Here, PINN w/ RAR & FD 20000 performs good

(They are superposing the ground-truth reference)



Experiment 3: Solving Inverse Problem

In experiment 3, the author wants to show that PINN can also be used to solve the inverse problem (Infer PDE parameters λ)

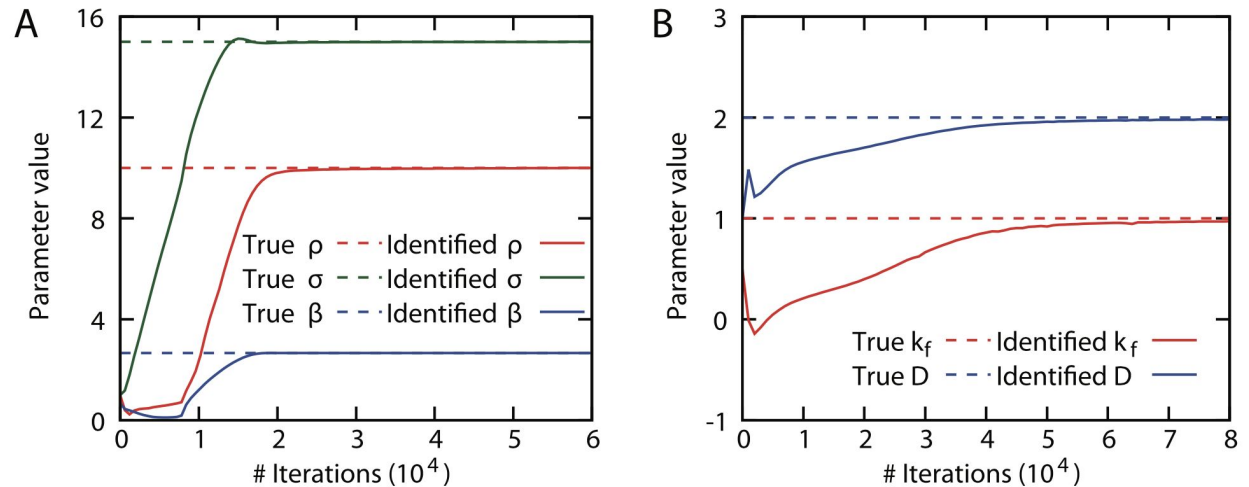


Fig. 9 Examples in subsections 4.3 and 4.4. The identified values of (A) the Lorenz system and (B) diffusion-reaction system converge to the true values during the training process. The parameter values are scaled for plotting.



Summary

Physics-Informed Neural Network (PINN) to embed a PDE into the loss of the neural network.

- When PDE parameters λ are known, PINN can learn NN parameter θ that approximates $u(x,t)$.
- When PDE parameters λ unknown but we have data information $I(u,x)=0$, PINN can learn both NN parameter θ and infer the PDE parameters λ .

Residual-based adaptive refinement (RAR) to improve the training efficiency.



Outline

1. DeepXDE: A Deep Learning Library of Solving Differential Equations. Lu L, Meng X, Mao Z, et al. SIAM Review 2021.
- 2. Solving high-dimensional partial differential equations using deep learning. Han J, Jentzen A, Weinan E. PNAS 2018.**
3. Forecasting Sequential Data Using Consistent Koopman Autoencoders. Azencot O, Erichson N B, Lin V, et al. ICML 2020.



Motivation & Idea

Motivation:

- Traditional methods to solve PDEs suffer from the curse of dimensionality.
- Deep learning methods can handle high dimensional PDEs.

Idea:

- Use neural networks to approximate the gradient of the unknown solution.

Mathematical setup

This paper focus on semilinear parabolic PDEs with the following form:

$$\begin{aligned} \frac{\partial u}{\partial t}(t, x) + \frac{1}{2} \text{Tr} \left(\sigma \sigma^{\text{T}}(t, x) (\text{Hess}_x u)(t, x) \right) + \nabla u(t, x) \cdot \mu(t, x) \\ + f \left(t, x, u(t, x), \sigma^{\text{T}}(t, x) \nabla u(t, x) \right) = 0 \end{aligned} \quad [1]$$

Where t is time, x is in d dimensions. μ and σ are functions.

Specifically, the author focus only on the solution at $t=0$, $x=\varepsilon$. Besides, X_t in the PDE can be stochastic process with Brownian motion:

$$X_t = \xi + \int_0^t \mu(s, X_s) ds + \int_0^t \sigma(s, X_s) dW_s.$$

Mathematical setup

Besides, we have two important equations:

$$X_{t_{n+1}} - X_{t_n} \approx \mu(t_n, X_{t_n}) \Delta t_n + \sigma(t_n, X_{t_n}) \Delta W_n, \quad [4]$$

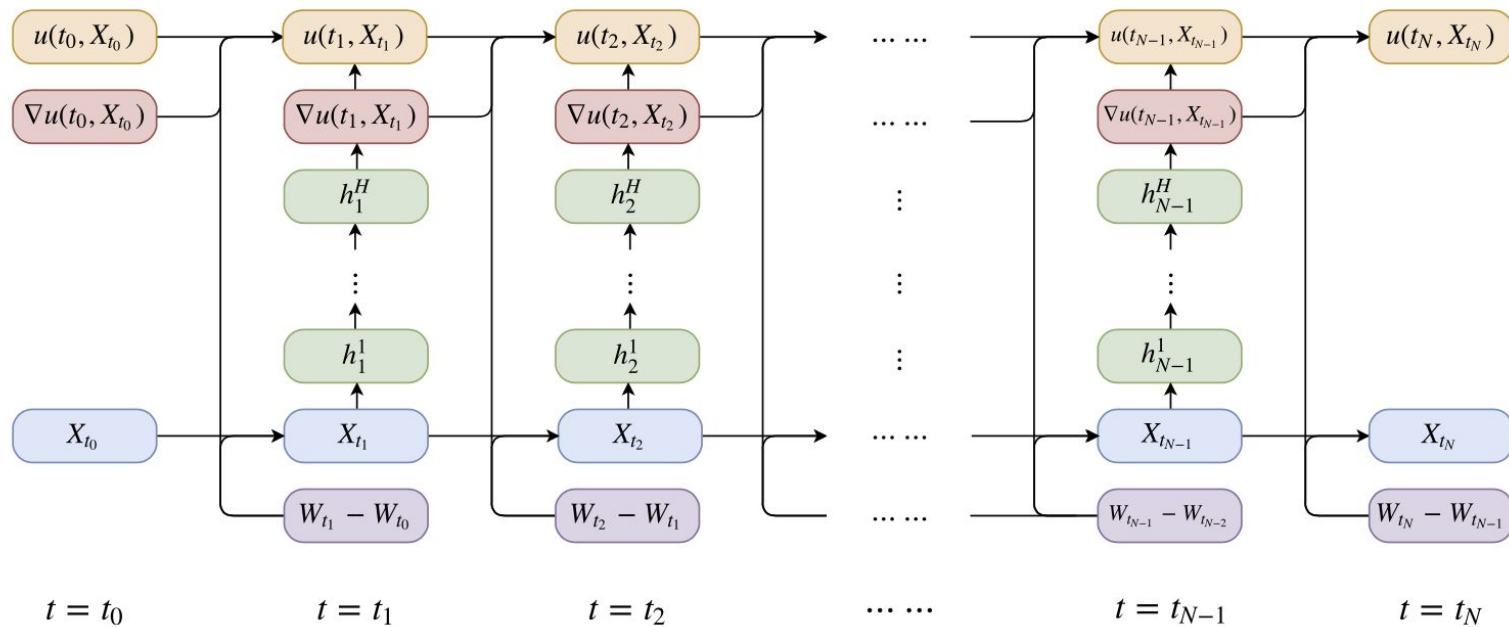
and

$$\begin{aligned} & u(t_{n+1}, X_{t_{n+1}}) - u(t_n, X_{t_n}) \\ & \approx -f\left(t_n, X_{t_n}, u(t_n, X_{t_n}), \sigma^\top(t_n, X_{t_n}) \nabla u(t_n, X_{t_n})\right) \Delta t_n \quad [5] \\ & + [\nabla u(t_n, X_{t_n})]^\top \sigma(t_n, X_{t_n}) \Delta W_n, \end{aligned}$$

This means given $\Delta u(t_n, X_{t_n})$, X_{t_n} (input) and W_{t_n} (input), the whole system will be known.

Therefore, the author wants to approximate $\Delta u(t_n, X_{t_n})$ using neural networks.

Framework

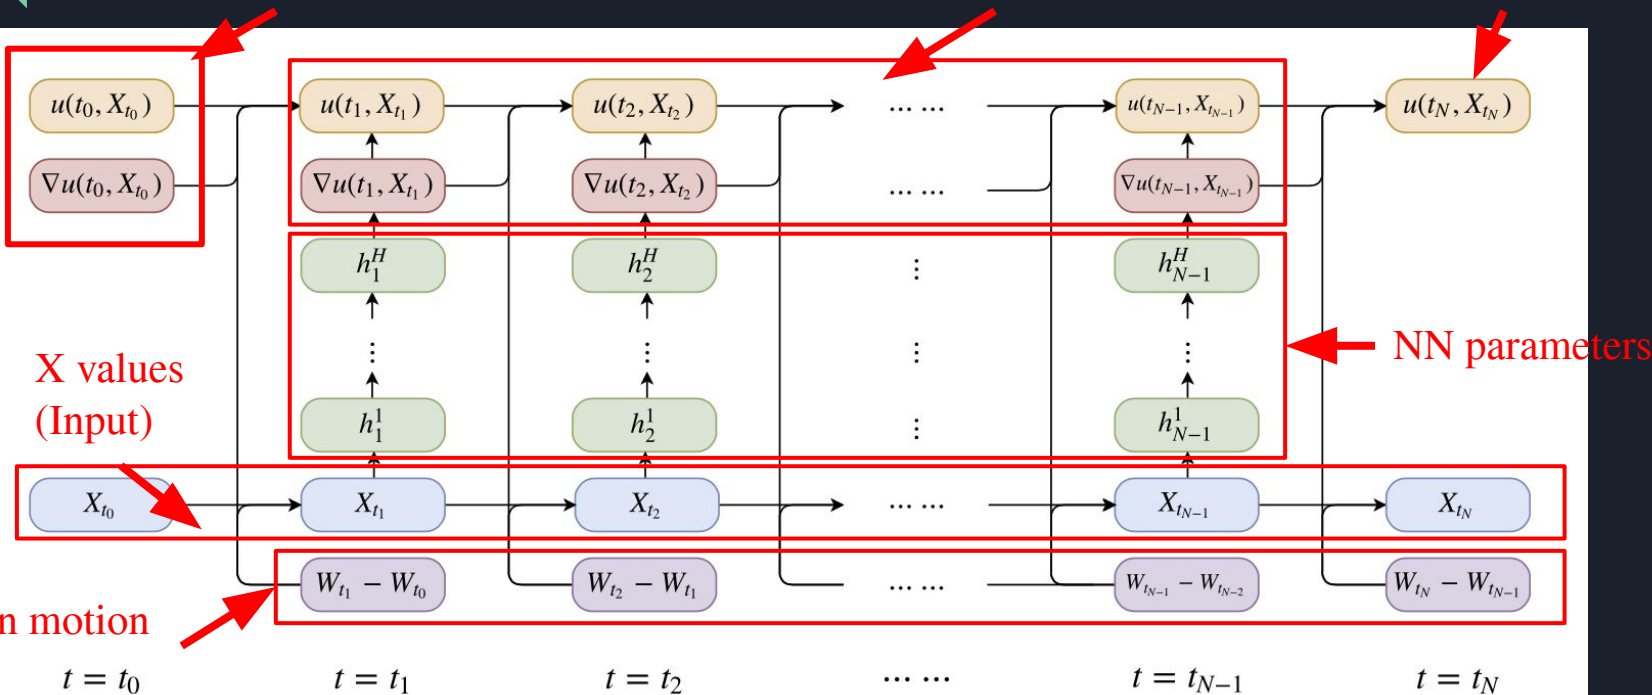


Framework

What we want to know

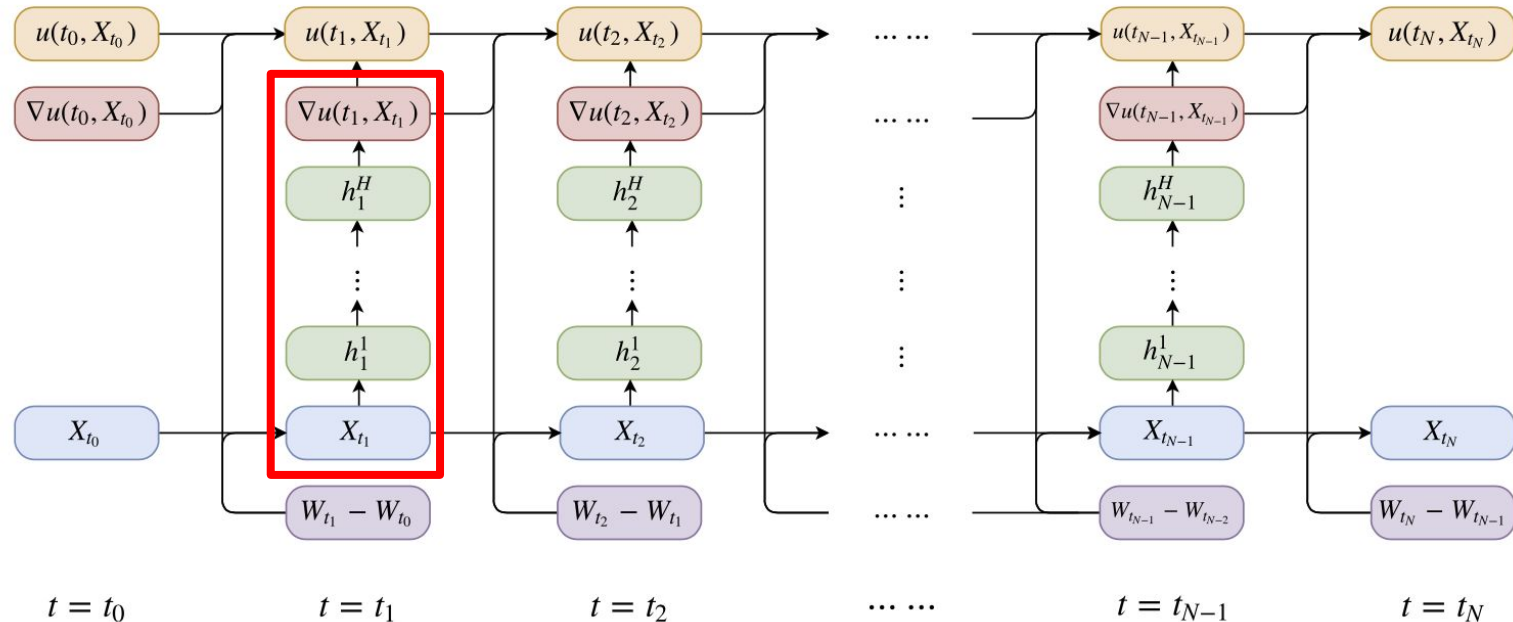
What we need to infer

What we use for loss



Framework

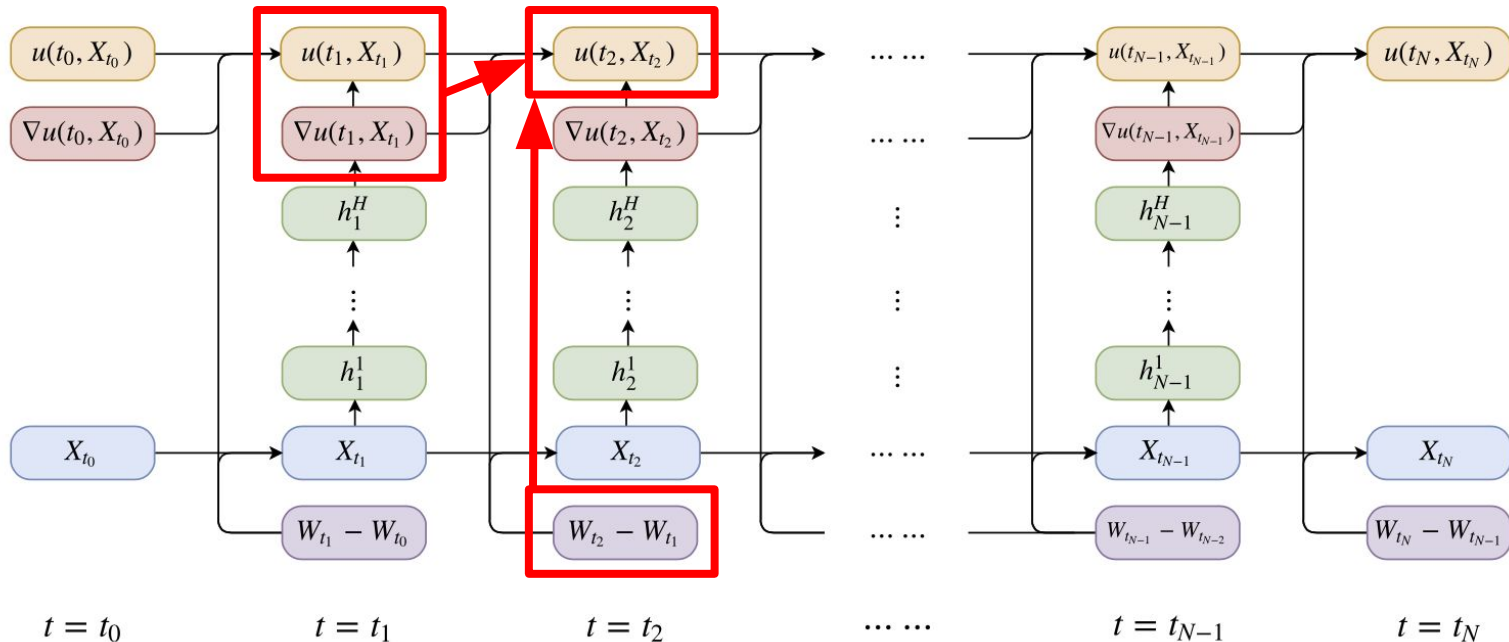
- (1) $X_{t_n} \rightarrow \text{del } u(t_n, X_{t_n})$: Where the NN parameters θ_n need to be learned.



Framework

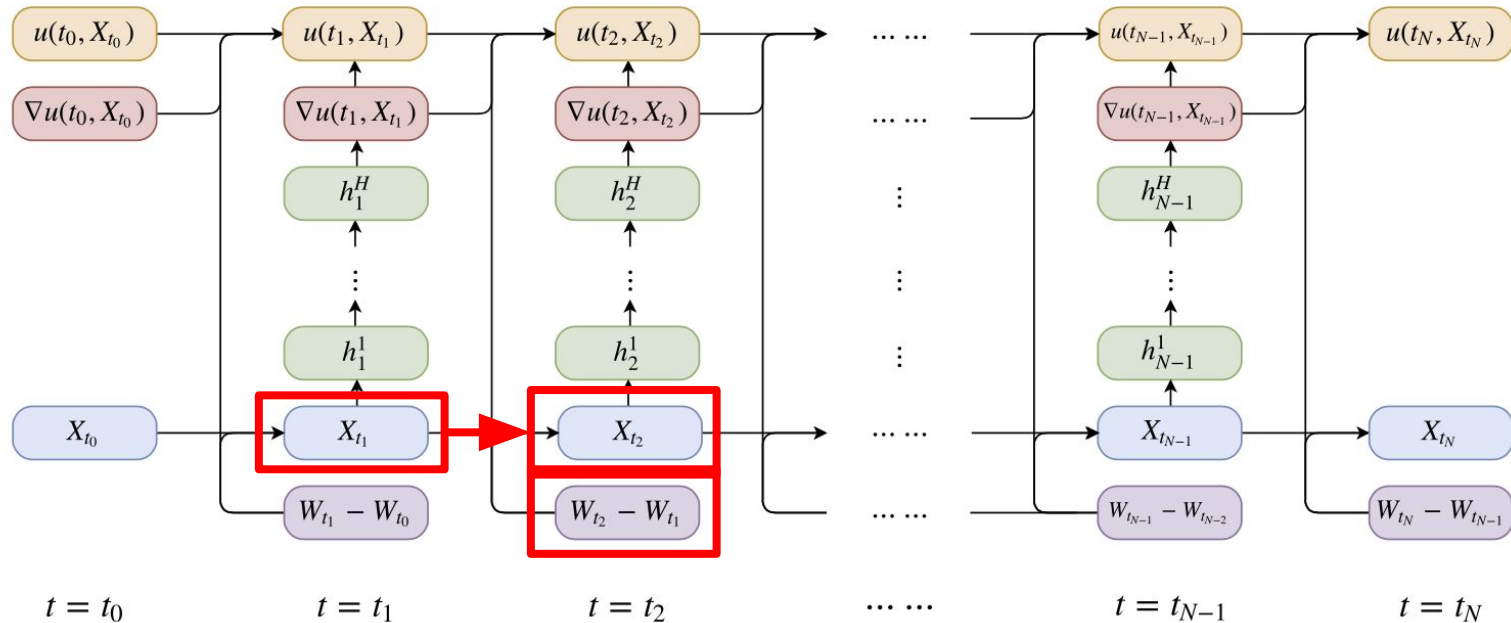
$$\begin{aligned}
 & u(t_{n+1}, X_{t_{n+1}}) - u(t_n, X_{t_n}) \\
 & \approx -f\left(t_n, X_{t_n}, u(t_n, X_{t_n}), \sigma^T(t_n, X_{t_n}) \nabla u(t_n, X_{t_n})\right) \Delta t_n \quad [5] \\
 & + [\nabla u(t_n, X_{t_n})]^T \sigma(t_n, X_{t_n}) \Delta W_n,
 \end{aligned}$$

(2) $u(t_n, X_{t_n}), \nabla u(t_n, X_{t_n}), W_{t_{n+1}} - W_{t_n} \rightarrow u(t_{n+1}, X_{t_{n+1}})$



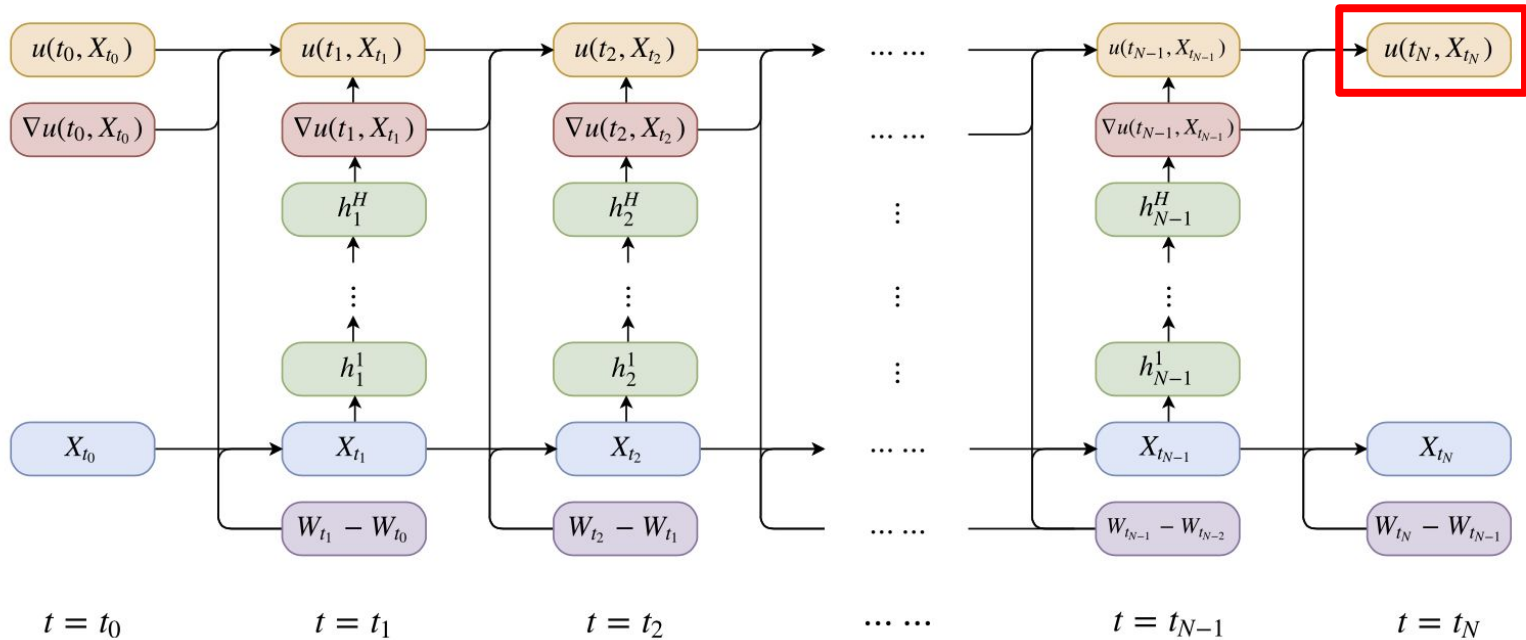
Framework

(3) $X_{t_n}, W_{t_{n+1}} - W_{t_n} \rightarrow X_{t_{n+1}}$ (X_t short cut)



Loss

(3) Loss = $\| u(t_N, X_{t_N}) - \text{estimated_}u(t_N, X_{t_N}) \|$





Outline

1. DeepXDE: A Deep Learning Library of Solving Differential Equations. Lu L, Meng X, Mao Z, et al. SIAM Review 2021.
2. Solving high-dimensional partial differential equations using deep learning. Han J, Jentzen A, Weinan E. PNAS 2018.
3. Forecasting Sequential Data Using Consistent Koopman Autoencoders. Azencot O, Erichson N B, Lin V, et al. ICML 2020.



Motivation & Idea

Motivation:

- Using RNN to forecast sequential data fault to capture the physical structures.
- Physics-based methods like Koopman autoencoder can help to capture such structures.

Idea:

- Use the Koopman autoencoder to analysis (and predict) the sequential data.



Mathematical Background

Given the sequential data: $Z_{k+1} = \varphi(Z_k)$, the Koopman operator satisfies $K\varphi f(z) = f(\varphi(z))$.

Similarly, we have $Z_{k-1} = \psi(Z_k)$ and $u\psi f(z) = f(\psi(z))$.

Besides, we have $C = XK\varphi X^{-1}$ and $D = Xu\psi X^{-1}$ where X is a finite transformation matrix.

Therefore, our work is to find C , D , and X .

To model X , the author use X_d to approximate X , X_e to approximate X^{-1}

Mathematical Background

Here, the author assume we are given the scalar observations of the system: $f_{k+1}(z) = f_k(\varphi(z))$.
Therefore, we have $f_{k+L} = f_k * \varphi^L$.

Besides, we will also have

- $f_{k+L} = X_d * C^L * X_e(f_k)$
- $f_{k-L} = X_d * D^L * X_e(f_k)$

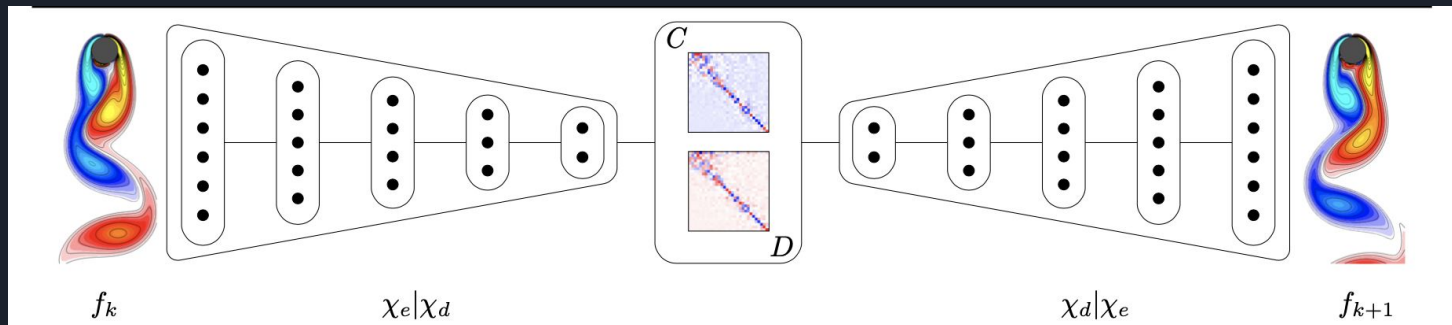


Figure 3. Our network takes observations f_k , and it learns a latent representation of them via an autoencoder architecture χ_e and χ_d . Then, the latent variables are propagated forward and backward in time using the linear layers C and D , respectively. We emphasize that all the above connections are bi-directional and so information can flow freely from left to right or right to left.

Loss 1: Identity Loss

For simplification, here we show the situation $L=1$.

- $f_{\{k+1\}} = \chi_d \circ \chi_e(f_k)$
- $f_{\{k-1\}} = \chi_d \circ \chi_e(f_k)$

The reconstructed version of f should be close to f

$$\tilde{f} = \chi_d \circ \chi_e(f),$$
$$\mathcal{E}_{\text{id}} = \frac{1}{2n} \sum_{k=1}^n \|f_k - \tilde{f}_k\|_2^2,$$

Where k is the time in range $[1, 2, \dots, n]$.

Loss 2 & 3: Forward and Backward Loss

- $f_{\{k+1\}} = Xd * C^1 * Xe(f_k)$
- $f_{\{k-1\}} = Xd * D^1 * Xe(f_k)$

The forecast f (forward and backward) should be close to ground truth.

$$\mathcal{E}_{\text{fwd}} = \frac{1}{2\lambda_s n} \sum_{l=1}^{\lambda_s} \sum_{k=1}^n \|f_{k+l} - \hat{f}_{k+l}\|_2^2,$$
$$\mathcal{E}_{\text{bwd}} = \frac{1}{2\lambda_s n} \sum_{l=1}^{\lambda_s} \sum_{k=1}^n \|f_{k-l} - \check{f}_{k-l}\|_2^2,$$

Here, we want to ensure our forecast to be well in λ_s steps forward and backward.

Loss 4: Forward and Backward Loss

- $f_{k+1} = Xd * C^{k+1} * X e(f_k)$
- $f_{k-1} = Xd * D^{k-1} * X e(f_k)$

$C * D$ should be close to identity matrix.

$$\mathcal{E}_{\text{con}} = \sum_{k=1}^{\kappa} \frac{1}{2k} \|D_{k*} C_{*k} - I_k\|_F^2 + \frac{1}{2k} \|C_{k*} D_{*k} - I_k\|_F^2 ,$$

Here, the D_{k*} are the upper k rows of D and C_{*k} are the leftmost k columns of C . I_k is k -dimension identity matrix. K is the dimension of C and D .