

Continuous-Discrete Differentiable Particle Filtering

Yunpeng Li


Major contributor: Xiongjie Chen

School of Computer Science and Electronic Engineering
University of Surrey, United Kingdom

Bellairs Workshop 2024

Motivating examples: Mars rover¹

- ▶ Continuous-time sequential state estimation.
- ▶ Missing observations, irregularly distributed time grid.



K9 Mars rover
Image source: NASA

¹Ng et al., "Continuous Time Particle Filtering", IJCAI, 2005.

Motivating examples: dental disease detection

- ▶ Continuous-time sequential state estimation.
- ▶ Missing observations, irregularly distributed time grid.



January 2019



February 2023

Some Initial Explorations of Differentiable Particle Filters

Yunpeng Li

Collaborators: Xiongjie Chen, Hao Wen, Ge
Conghui Hu

Department of Computer Science
University of Surrey, UK


Bellairs Workshop 2021

Future directions

- ▶ Online learning.
- ▶ Better dynamic models – e.g. mixture models?
- ▶ Better measurement models.
- ▶ Better proposal distributions – e.g. APF with normalizing flow, physics-inspired NNs?
- ▶ Better optimisation objectives – existing VI-based, self-supervised methods?
- ▶ Differentiable resampling schemes¹³.
- ▶ Distributed learning and inference.
- ▶ Continuous-time filtering.

¹³ Corenflos et al., "Differentiable Particle Filtering via Entropy-Regularized Optimal Transport", ICML, 2021.

Continuous-discrete state-space models



- ▶ Dynamic model:

$$ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$$

- ▶ Measurement model:

$$o_{t_k} = H_{t_k}(s_{t_k}, v_{t_k}, \theta)$$

s_{t_k} the hidden state at time t_k

θ model parameters


$f_\theta, \sigma_\theta, H_t$ deterministic functions

o_{t_k} the observation at time t_k

dB_t Brownian motion

v_{t_k} measurement noise


Discrete state-space models



- ▶ Dynamic model:
 $s_t = K_t(s_{t-1}, u_t, \theta)$
 s_t the hidden state at time t
 θ model parameters
- ▶ Measurement model:
 $o_t = H_t(s_t, v_t, \theta)$
 f_θ, K_t, H_t deterministic functions
 o_t the observation at time t
 u_t, v_t noise terms

Examples of continuous-discrete state-space models


Continuous-discrete stochastic volatility model



- ▶ Dynamic model:
 $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$
 $ds_t = (\eta - 1)(s_t - \mu)dt + \beta dB_t$
- ▶ Measurement model:
 $o_{t_k} = H_{t_k}(s_{t_k}, v_{t_k}, \theta)$
 $o_{t_k} | s_{t_k} \sim \mathcal{N}(0, \gamma^2 \exp(s_t))$

Examples of continuous-discrete state-space models

Discrete stochastic volatility model: evenly spaced time instances.

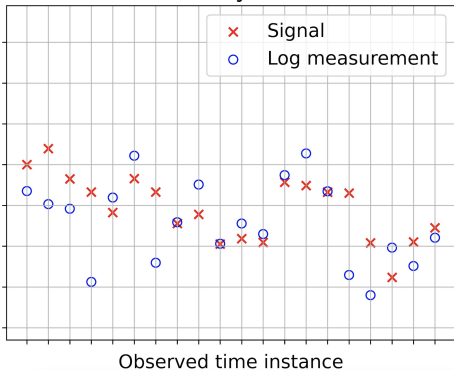


- ▶ Dynamic model:
 $s_t = K_t(s_{t-1}, u_t, \theta)$
- ▶ Measurement model:
 $o_t = H_t(s_t, v_t, \theta)$
- ▶ $s_t = \eta(s_{t-1} - \mu) + \mu + u_t$,
 $u_t \sim \mathcal{N}(0, \beta^2)$
- ▶ $o_t | s_t \sim \mathcal{N}(0, \gamma^2 \exp(s_t))$

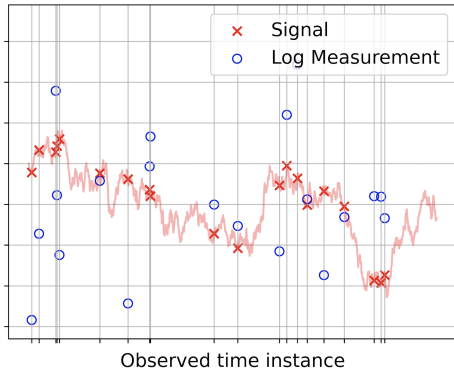
Examples of state-space models

Comparison of discrete and continuous-discrete state-space models

Stochastic volatility model




Discrete



Continuous-discrete

Examples of continuous-discrete state-space models

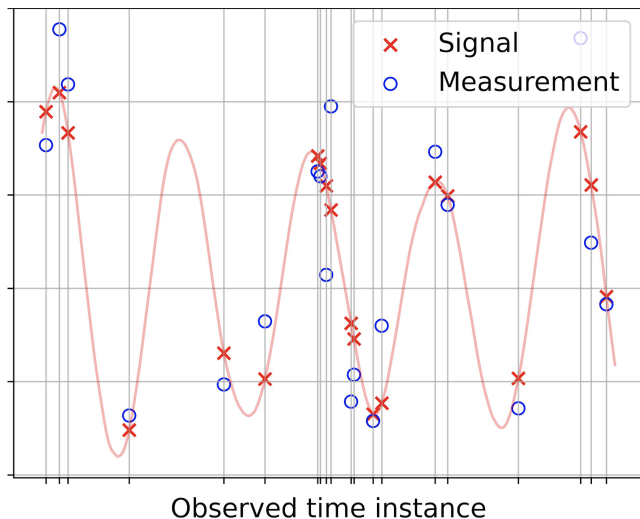
Noisy pendulum model



- ▶ Dynamic model:
 $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$
 $ds_t(1) = s_t(2)dt$
 $ds_t(2) = -a^2 \sin(s_t(1))dt + b^{\frac{1}{2}}dB_t$
- ▶ Measurement model:
 $o_{t_k} = H_{t_k}(s_{t_k}, v_{t_k}, \theta)$
 $o_{t_k} | s_{t_k} \sim \mathcal{N}(s_t(1), \sigma^2)$


Examples of state-space models

Noisy pendulum model



Continuous-discrete filtering problem formulation

Continuous-discrete filtering




- ▶ Recursively estimate the posterior distribution of latent states $p(s_{t_{1:k}} | o_{t_{1:k}}; \theta)$ in continuous-time.
- ▶ We can infer the posterior at arbitrary time instances.

¹Ng et al., "Continuous Time Particle Filtering", IJCAI, 2005.

²Bucy et al., "Filtering for Stochastic Processes with Applications to Guidance", American Mathematical Society, 2005.

Continuous-discrete filtering problem formulation

Continuous-discrete filtering



- ▶ Recursively estimate the posterior distribution of latent states $p(s_{t_{1:k}} | o_{t_{1:k}}; \theta)$ in continuous-time.
- ▶ We can infer the posterior at arbitrary time instances.


Linear Gaussian models: [Kalman-Bucy filters](#)².

¹Ng et al., "Continuous Time Particle Filtering", IJCAI, 2005.

²Bucy et al., "Filtering for Stochastic Processes with Applications to Guidance", American Mathematical Society, 2005.

Continuous-discrete filtering problem formulation

Continuous-discrete filtering



- ▶ Recursively estimate the posterior distribution of latent states $p(s_{t_{1:k}} | o_{t_{1:k}}; \theta)$ in continuous-time.
- ▶ We can infer the posterior at arbitrary time instances.

Linear Gaussian models: [Kalman-Bucy filters](#)².

Non-linear non-Gaussian models: [Continuous-discrete particle filters](#)¹.

¹Ng et al., "Continuous Time Particle Filtering", IJCAI, 2005.

²Bucy et al., "Filtering for Stochastic Processes with Applications to Guidance", American Mathematical Society, 2005.

Components

- ▶ System dynamics and proposal process are specified by stochastic differential equations (SDEs).
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are deterministic functions.

Components

- ▶ System dynamics and proposal process are specified by stochastic differential equations (SDEs).
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are deterministic functions.

- ▶ Measurement model: $p(o_{t_k} | s_{t_k}^i; \theta)$.

Continuous-discrete particle filtering

Components

- ▶ System dynamics and proposal process are specified by stochastic differential equations (SDEs).

- ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

- ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are deterministic functions.

- ▶ Measurement model: $p(o_{t_k} | s_{t_k}^i; \theta)$.

- ▶ Update particles weights: $w_{t_k}^i = w_{t_{k-1}}^i \frac{p(o_{t_k} | s_{t_k}^i; \theta)p(s_{t_k}^i | s_{t_{k-1}}^i; \theta)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k}; \phi)}$.

Continuous-discrete particle filtering

Components

- ▶ System dynamics and proposal process are specified by stochastic differential equations (SDEs).

- ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

- ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are deterministic functions.

- ▶ Measurement model: $p(o_{t_k} | s_{t_k}^i; \theta)$.

- ▶ Update particles weights: $w_{t_k}^i = w_{t_{k-1}}^i \frac{p(o_{t_k} | s_{t_k}^i; \theta)p(s_{t_k}^i | s_{t_{k-1}}^i; \theta)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k}; \phi)}$.

- ▶ Resampling.

Continuous-discrete particle filtering

Initialisation: Draw $\{s_{t_0}^i\}_{i=1}^{N_p}$ from $p(s_{t_0})$.

Set $\{\tilde{w}_{t_0}^i = \frac{1}{N_p}\}_{i=1}^{N_p}$.

for $k = 1$ to K :

for $i = 1$ to N_p :

$$s_{t_k}^i = \int_{t_{k-1}}^{t_k} g_\phi(s_t, o_t, t) dt + \int_{t_{k-1}}^{t_k} \sigma_\theta(s_t, t) dB_t \text{ with } s_{t_{k-1}} = s_{t_{k-1}}^i.$$

$$\text{Update weights: } w_{t_k}^i = w_{t_{k-1}}^i \frac{p(o_{t_k} | s_{t_k}^i; \theta) p(s_{t_k}^i | s_{t_{k-1}}^i; \theta)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k}; \phi)}.$$

end for

for $i = 1$ to N_p :

$$\text{Normalise weights: } \tilde{w}_{t_k}^i = w_{t_k}^i / \sum_{j=1}^{N_p} w_{t_k}^j.$$

end for

if ESS < threshold:

$$\text{Resample } \{s_{t_k}^i, \tilde{w}_{t_k}^i\}_{i=1}^{N_p} \text{ to obtain } \{s_{t_k}^i, \frac{1}{N_p}\}_{i=1}^{N_p}.$$

end if

end for

Construct continuous-discrete particle filters (DPFs) with machine learning tools:

- ▶ Build system dynamics and proposal processes with neural stochastic differential equations (neural SDEs)^{3, 4, 5}.
- ▶ Build measurement models with neural networks⁶.

³Li et al., “Scalable gradients for stochastic differential equations”, AISTATS, 2020.

⁴Deng et al., “Continuous Latent Process Flows”, NeurIPS, 2021.

⁵Deng et al., “Continuous-time Particle Filtering for Latent Stochastic Differential Equations”, arXiv, 2209.00173, 2022.

⁶Chen and Li, “Conditional Measurement Density Estimation in Sequential Monte Carlo via Normalizing Flow”, EUSIPCO, 2022.

Neural ordinary differential equations

Neural ordinary differential equations (Neural ODEs)⁷:

$$\frac{ds_t}{dt} = f_{\theta}(s_t, t), s_0 = s(0),$$

model the dynamic function $f_{\theta}(s_t, t)$ with neural networks.

⁷Chen et al., “Neural ordinary differential equations”, NeurIPS, 2018.

Neural ordinary differential equations

Neural ordinary differential equations (Neural ODEs)⁷:

$$\frac{ds_t}{dt} = f_{\theta}(s_t, t), s_0 = s(0),$$

model the dynamic function $f_{\theta}(s_t, t)$ with neural networks.

- ▶ How to backpropagate gradients through ODE solvers?
 - ▶ Adjoint sensitivity method.
 - ▶ Can be trained in the same way as normal neural networks.

- ▶ Applications:
 - ▶ Irregularly-sampled time series modelling.
 - ▶ Continuous normalising flows.

⁷Chen et al., “Neural ordinary differential equations”, NeurIPS, 2018.

Neural stochastic differential equations

Neural stochastic differential equations (Neural SDEs)⁸:

$$ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t,$$

a stochastic variant of neural ODEs.

- ▶ Li et al. extended the adjoint sensitivity method developed for neural ODEs to neural SDEs³ - we can backpropagate through SDE solvers.
- ▶ A natural choice when dealing with stochastic dynamic systems.

³Li et al., "Scalable gradients for stochastic differential equations", AISTATS, 2020.

⁸Tzen and Raginsky, "Neural stochastic differential equations: Deep latent Gaussian models in the diffusion limit", arXiv, 1905.09883, 2019.

- ▶ Dynamic model.
 - ▶ Construct with neural SDEs.

$$ds_{t_k} = f_{\theta}(s_t, t)dt + \sigma_{\theta}(s_t, t)dB_t,$$
$$s_{t_k} = \int_{t_{k-1}}^{t_k} f_{\theta}(s_t, t)dt + \int_{t_{k-1}}^{t_k} \sigma_{\theta}(s_t, t)dB_t.$$

Build system dynamics and proposals with neural SDEs

- ▶ Dynamic model.
 - ▶ Construct with neural SDEs.

$$ds_{t_k} = f_{\theta}(s_t, t)dt + \sigma_{\theta}(s_t, t)dB_t,$$
$$s_{t_k} = \int_{t_{k-1}}^{t_k} f_{\theta}(s_t, t)dt + \int_{t_{k-1}}^{t_k} \sigma_{\theta}(s_t, t)dB_t.$$

- ▶ Proposal process.
 - ▶ Construct with neural SDEs, include observations as neural network inputs.

$$ds_{t_k} = g_{\phi}(s_t, \mathbf{o}_t, t)dt + \sigma_{\theta}(s_t, t)dB_t,$$
$$s_{t_k} = \int_{t_{k-1}}^{t_k} g_{\phi}(s_t, \mathbf{o}_t, t)dt + \int_{t_{k-1}}^{t_k} \sigma_{\theta}(s_t, t)dB_t.$$

- ▶ Latent SDEs³, continuous latent process flow (CLPF)⁴.
 - ▶ Generate variational posterior distributions.
 - ▶ Model observations as a continuous stochastic process.
- ▶ Continuous-time particle filters (CTPF)⁵.
 - ▶ Non-differentiable resampling.

³Li et al., “Scalable gradients for stochastic differential equations”, AISTATS, 2020.

⁴Deng et al., “Continuous Latent Process Flows”, NeurIPS, 2021.

⁵Deng et al., “Continuous-time Particle Filtering for Latent Stochastic Differential Equations”, arxiv, 2209.00173, 2022.

Updating particle weights can be difficult:

1. Intractable transition densities.
2. How to design flexible measurement models?

Continuous-discrete DPF: weights update

Recall that when updating particle weights:

$$w_{t_k}^i = w_{t_{k-1}}^i \frac{p(s_{t_k}^i | s_{t_{k-1}}^i) p(o_{t_k} | s_{t_k}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

Continuous-discrete DPF: weights update

Recall that when updating particle weights:

$$w_{t_k}^i = w_{t_{k-1}}^i \frac{p(s_{t_k}^i | s_{t_{k-1}}^i) p(o_{t_k} | s_{t_k}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

Transition density ratio $\frac{p(s_{t_k}^i | s_{t_{k-1}}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$.

- ▶ Neither $p(s_{t_k}^i | s_{t_{k-1}}^i)$ nor $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})$ are tractable in continuous-discrete state-space models, because $p(s_{t_k}^i | s_{t_{k-1}}^i)$ and $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})$ are implicitly defined by neural SDEs.

Continuous-discrete DPF: weights update

How to update particle weights?

$$w_{t_k}^i = w_{t_{k-1}}^i \frac{p(s_{t_k}^i | s_{t_{k-1}}^i) p(o_{t_k} | s_{t_k}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

- ▶ Restrict to bootstrap filtering approaches, so that $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})$ and $p(s_{t_k}^i | s_{t_{k-1}}^i)$ are cancelled.

Continuous-discrete DPF: weights update

How to update particle weights?

$$w_{t_k}^i = w_{t_{k-1}}^i \frac{p(s_{t_k}^i | s_{t_{k-1}}^i) p(o_{t_k} | s_{t_k}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

- ▶ Restrict to bootstrap filtering approaches, so that $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})$ and $p(s_{t_k}^i | s_{t_{k-1}}^i)$ are cancelled.
 - ▶ No, this will lead to very low sampling efficiency and high variance estimators.

Continuous-discrete DPF: weights update

How to update particle weights?

$$w_{t_k}^i = w_{t_{k-1}}^i \frac{p(s_{t_k}^i | s_{t_{k-1}}^i) p(o_{t_k} | s_{t_k}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

- ▶ Restrict to bootstrap filtering approaches, so that $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})$ and $p(s_{t_k}^i | s_{t_{k-1}}^i)$ are cancelled.
 - ▶ No, this will lead to very low sampling efficiency and high variance estimators.
- ▶ Random weight approaches.
 - ▶ Update particle weights with unbiased estimators of the ratio

$$\frac{p(s_{t_k}^i | s_{t_{k-1}}^i)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k})}$$

Continuous-discrete DPF: weights update

- ▶ Importance sampling: $\mathbb{E}_p[f(x)] = \mathbb{E}_q\left[\frac{p(x)}{q(x)}f(x)\right]$.

How to estimate $\mathbb{E}_p[f(x)]$?

- ▶ Sample from q .
- ▶ Estimate $\mathbb{E}_p[f(x)] \approx \frac{1}{N} \sum_{i=1}^N \frac{p(x_i)}{q(x_i)} f(x_i)$.

- ▶ Random weight importance sampling⁹.

- ▶ Unbiased estimator of the ratio: $\mathbb{E}[Q] = \frac{p(x)}{q(x)}$.

- ▶ $\mathbb{E}_p[f(x)] = \mathbb{E}_q\left[\frac{p(x)}{q(x)}f(x)\right] = \mathbb{E}_q\left[\mathbb{E}[Q]f(x)\right]$.

How to estimate $\mathbb{E}_p[f(x)]$?

- ▶ Sample from q .
- ▶ Draw samples of Q .
- ▶ Estimate $\mathbb{E}_p[f(x)] \approx \frac{1}{N} \sum_{i=1}^N \frac{1}{M} \sum_{j=1}^M Q_j f(x_i)$.

⁹Chopin and Papaspiliopoulos, "An Introduction to Sequential Monte Carlo", Springer, 2020.

Continuous-discrete DPF: weights update

- ▶ Importance sampling: $\mathbb{E}_p[f(x)] = \mathbb{E}_q\left[\frac{p(x)}{q(x)}f(x)\right]$.

How to estimate $\mathbb{E}_p[f(x)]$?

- ▶ Sample from q .
- ▶ Estimate $\mathbb{E}_p[f(x)] \approx \frac{1}{N} \sum_{i=1}^N \frac{p(x_i)}{q(x_i)} f(x_i)$.

- ▶ Random weight importance sampling⁹.

- ▶ Unbiased estimator of the ratio: $\mathbb{E}[Q] = \frac{p(x)}{q(x)}$.

- ▶ $\mathbb{E}_p[f(x)] = \mathbb{E}_q\left[\frac{p(x)}{q(x)}f(x)\right] = \mathbb{E}_q\left[\mathbb{E}[Q]f(x)\right]$.

How to estimate $\mathbb{E}_p[f(x)]$?

- ▶ Sample from q .
- ▶ Draw samples of Q .
- ▶ Estimate $\mathbb{E}_p[f(x)] \approx \frac{1}{N} \sum_{i=1}^N \frac{1}{M} \sum_{j=1}^M Q_j f(x_i)$.

⁹Chopin and Papaspiliopoulos, "An Introduction to Sequential Monte Carlo", Springer, 2020.

Continuous-discrete DPF: weights update

Random weight particle filtering^{10,11}.

Initialisation: Draw $\{s_{t_0}^i\}_{i=1}^{N_p}$ from $p(s_{t_0})$.

Set $\{\tilde{w}_{t_0}^i = \frac{1}{N_p}\}_{i=1}^{N_p}$.

for $k = 1$ to K :

for $i = 1$ to N_p :

Draw $s_{t_k}^i$ from $q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k}; \phi)$.

Draw samples $\{Q_{t_k}^j\}_{j=1}^M$, with $\mathbb{E}[Q] = \frac{p(o_{t_k} | s_{t_k}^i; \theta) p(s_{t_k}^i | s_{t_{k-1}}^i; \theta)}{q(s_{t_k}^i | s_{t_{k-1}}^i, o_{t_k}; \phi)}$.

Update weights: $w_{t_k}^i = w_{t_{k-1}}^i \frac{1}{M} \sum_{j=1}^M Q_j$.

end for

for $i = 1$ to N_p :

Normalise weights: $\tilde{w}_{t_k}^i = w_{t_k}^i / \sum_{j=1}^{N_p} w_{t_k}^j$.

end for

if ESS < threshold:

Resample $\{s_{t_k}^i, \tilde{w}_{t_k}^i\}_{i=1}^{N_p}$ to obtain $\{s_{t_k}^i, \frac{1}{N_p}\}_{i=1}^{N_p}$.

end if

end for

¹⁰Fearnhead et al., "Random-weight particle filtering of continuous time processes", JRSSB, 2010.

¹¹Fearnhead et al., "Particle filters for partially observed diffusions", JRSSB, 2008.

Continuous-discrete DPFs: weights update

- ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
- ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

An unbiased estimator of $\frac{p(s_{t_{k+1}}^i | s_{t_k}^i; \theta)}{q(s_{t_{k+1}}^i | s_{t_k}^i, o_{t_1:k}; \phi)}$ derived from the Girsanov theorem:

$$\begin{aligned} & Z(t_k, t_{k+1}; \omega_{t_{k+1}}^i) \\ &= \exp \left(\int_{t_k}^{t_{k+1}} \left[f_\theta(s_t, t) - g_\phi(s_t, o_{t_k}, t) \right]^\top \left[\sigma_\theta^{-1}(s_t, t) \right]^\top dB_{t_{k+1}-t_k}^i \right. \\ & \quad \left. - \frac{1}{2} \int_{t_k}^{t_{k+1}} \left[f_\theta(s_t, t) - g_\phi(s_t, o_{t_k}, t) \right]^\top \left[\sigma_\theta(s_t, t) \sigma_\theta^\top(s_t, t) \right]^{-1} \right. \\ & \quad \left. \left[f_\theta(s_t, t) - g_\phi(s_t, h_{t_k}, t) \right] dt \right), \\ &= \exp \left(\int_{t_k}^{t_{k+1}} F_{\theta, \phi}(s_t, o_{t_k}, t) dt + \int_{t_k}^{t_{k+1}} G_{\theta, \phi}(s_t, o_{t_k}, t) dB_{t_{k+1}-t_k}^i \right). \end{aligned}$$

How to compute the Ito integral $Z(t_k, t_{k+1}; \omega_{k+1}^i)$?

- ▶ Augment the latent state dimension with an concatenated dimension.

$$dS_t = \begin{bmatrix} ds_t \\ ds'_t \end{bmatrix} = \begin{bmatrix} g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t \\ F_{\theta, \phi}(s_t, o_{t_k}, t)dt + G_{\theta, \phi}(s_t, o_{t_k}, t)dB_t \end{bmatrix} \quad (1)$$

¹²Higham, "An algorithmic introduction to numerical simulation of stochastic differential equations", SIAM review, 2001.

How to compute the Ito integral $Z(t_k, t_{k+1}; \omega_{k+1}^i)$?

- ▶ Augment the latent state dimension with an concatenated dimension.

$$dS_t = \begin{bmatrix} ds_t \\ ds'_t \end{bmatrix} = \begin{bmatrix} g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t \\ F_{\theta, \phi}(s_t, o_{t_k}, t)dt + G_{\theta, \phi}(s_t, o_{t_k}, t)dB_t \end{bmatrix} \quad (1)$$

- ▶ Solve the concatenated SDE using SDE solvers, e.g. Euler-Maruyama and Runge-Kutta methods¹².
 - ▶ Simultaneously draw particles and compute $Z(t_k, t_{k+1}; \omega_{k+1}^i)$ by solving the concatenated SDE.

¹²Higham, "An algorithmic introduction to numerical simulation of stochastic differential equations", SIAM review, 2001.

Continuous-discrete DPF: weights update

Continuous-discrete DPF

Initialisation: Draw $\{s_{t_0}^i\}_{i=1}^{N_p}$ from $p(s_{t_0})$.

Set $\{\tilde{w}_{t_0}^i = \frac{1}{N_p}\}_{i=1}^{N_p}$.

for $k = 1$ to K :

for $i = 1$ to N_p :

Draw $s_{t_k}^i$ and estimate $Z(t_{k-1}, t_k; \omega_k^i)$ by solving Equation (1).

Update weights: $w_{t_k}^i = w_{t_{k-1}}^i p(o_{t_k} | s_{t_k}^i; \theta) Z(t_{k-1}, t_k; \omega_k^i)$.

end for

for $i = 1$ to N_p :

Normalise weights: $\tilde{w}_{t_k}^i = w_{t_k}^i / \sum_{j=1}^{N_p} w_{t_k}^j$.

end for

if ESS < threshold:

Resample $\{s_{t_k}^i, \tilde{w}_{t_k}^i\}_{i=1}^{N_p}$ to obtain $\{s_{t_k}^i, \frac{1}{N_p}\}_{i=1}^{N_p}$.

end if

end for


Theorem

Given $s_{t_{k+1}}^i$, $s_{t_k}^i$, and $\omega_{k+1}^i \in \Omega_{k+1}$ an outcome that generates a standard Brownian motion $B_{t_{k+1}-t_k}^i$ driving $s_{t_k}^i$ to $s_{t_{k+1}}^i$ in the proposal process, $Z(t_k, t_{k+1}; \omega_{k+1}^i)$ is an unbiased estimator of the transition density ratio

$$\frac{p(s_{t_{k+1}}^i | s_{t_k}^i; \theta)}{q(s_{t_{k+1}}^i | s_{t_k}^i, o_{t_{1:k}}; \phi)}.$$

Continuous-discrete DPFs: model design

Conditional Normalising Flow-based Measurement Model⁶.




1. Map the observation o_t to a variable y_t^i that follows a Gaussian distribution $p_Y(\cdot)$ through the conditional normalising flow \bar{G}_θ .

⁶Chen et al., "Conditional Measurement Density Estimation in Sequential Monte Carlo via Normalising Flow", EUSIPCO, 2022.

Continuous-discrete DPFs: model design

Conditional Normalising Flow-based Measurement Model⁶.



1. Map the observation o_t to a variable y_t^i that follows a Gaussian distribution $p_Y(\cdot)$ through the conditional normalising flow $\bar{\mathcal{G}}_\theta$.
2. Compute the likelihood of observations o_t given state s_t using the change of variable formula: $p(o_t | s_t^i; \theta) = p_Y(\bar{\mathcal{G}}_\theta(o_t, s_t^i)) \left| \det \frac{\partial \bar{\mathcal{G}}_\theta(o_t, s_t^i)}{\partial o_t} \right|$.

⁶Chen et al., "Conditional Measurement Density Estimation in Sequential Monte Carlo via Normalising Flow", EUSIPCO, 2022.

An overview

- ▶ Construct dynamic process and proposal process with neural SDEs.
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
- f_θ , g_ϕ , and σ_θ are neural networks.

¹³Corenflos et al. "Differentiable Particle Filtering via Entropy-regularized Optimal Transport." ICML, 2021.

Continuous-discrete DPFs: model design

An overview

- ▶ Construct dynamic process and proposal process with neural SDEs.
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are neural networks.
- ▶ Conditional normalising flow-based measurement model.
 - ▶ $p(o_t | s_t^i; \theta) = p_Y(\bar{\mathcal{G}}_\theta(o_t, s_t^i)) \left| \det \frac{\partial \bar{\mathcal{G}}_\theta(o_t, s_t^i)}{\partial o_t} \right|$.

¹³Corenflos et al. "Differentiable Particle Filtering via Entropy-regularized Optimal Transport." ICML, 2021.

Continuous-discrete DPFs: model design

An overview

- ▶ Construct dynamic process and proposal process with neural SDEs.
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are neural networks.
- ▶ Conditional normalising flow-based measurement model.
 - ▶ $p(o_t | s_t^i; \theta) = p_Y(\bar{\mathcal{G}}_\theta(o_t, s_t^i)) \left| \det \frac{\partial \bar{\mathcal{G}}_\theta(o_t, s_t^i)}{\partial o_t} \right|$.
- ▶ Update weights with the unbiased estimator $Z(t_k, t_{k+1}; \omega_{k+1}^i)$.
 - ▶ $w_{t_{k+1}}^i \propto w_{t_k}^i p(o_{t_{k+1}} | s_{t_{k+1}}^i; \theta) Z(t_k, t_{k+1}; \omega_{k+1}^i)$.

¹³Corenflos et al. "Differentiable Particle Filtering via Entropy-regularized Optimal Transport." ICML, 2021.

Continuous-discrete DPFs: model design

An overview

- ▶ Construct dynamic process and proposal process with neural SDEs.
 - ▶ Dynamic process: $ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t$.
 - ▶ Proposal process: $ds_t = g_\phi(s_t, o_t, t)dt + \sigma_\theta(s_t, t)dB_t$.

f_θ , g_ϕ , and σ_θ are neural networks.
- ▶ Conditional normalising flow-based measurement model.
 - ▶ $p(o_t | s_t^i; \theta) = p_Y(\bar{\mathcal{G}}_\theta(o_t, s_t^i)) \left| \det \frac{\partial \bar{\mathcal{G}}_\theta(o_t, s_t^i)}{\partial o_t} \right|$.
- ▶ Update weights with the unbiased estimator $Z(t_k, t_{k+1}; \omega_{k+1}^i)$.
 - ▶ $w_{t_{k+1}}^i \propto w_{t_k}^i p(o_{t_{k+1}} | s_{t_{k+1}}^i; \theta) Z(t_k, t_{k+1}; \omega_{k+1}^i)$.
- ▶ Differentiable resampling.
 - ▶ Entropy-regularised optimal transport resampling¹³.

¹³Corenflos et al. "Differentiable Particle Filtering via Entropy-regularized Optimal Transport." ICML, 2021.

Continuous-discrete DPF: training objective

End-to-End learning by minimising a given loss function:

1. Supervised loss;

- ▶ The mean squared error (MSE):

$$L_{MSE}(\theta, \phi) = \frac{1}{K} \sum_{k=1}^K (s_{t_k}^* - s_{t_k})^T (s_{t_k}^* - s_{t_k}),$$

$s_{t_k}^*$: ground truth state, s_{t_k} : estimated states.

2. Unsupervised loss.

- ▶ SMC evidence lower bound:

$$L_{ELBO}(\theta, \phi) = \mathbb{E} \left[\log \hat{p}_{N_p}(o_{t_{1:K}}; \theta) \right],$$

$\hat{p}(o_{1:t}; \theta)$: an estimate of the marginal likelihood computed with particle weights.

Continuous-discrete DPF: filtering

A set of observed time instances $\{t_1, \dots, t_K\}$ with observations $\{o_{t_1}, \dots, o_{t_K}\}$.

Initialisation: Draw $\{s_{t_0}^i\}_{i=1}^{N_p}$ from $p(s_{t_0})$.

Set $\{\tilde{w}_{t_0}^i = \frac{1}{N_p}\}_{i=1}^{N_p}$.

for $k = 1$ to K :

for $i = 1$ to N_p :

Draw $s_{t_k}^i$ and estimate $Z(t_{k-1}, t_k; \omega_k^i)$ by solving Equation (1).

Update weights: $w_{t_k}^i = w_{t_{k-1}}^i p(o_{t_k} | s_{t_k}^i; \theta) Z(t_{k-1}, t_k; \omega_k^i)$.

end for

for $i = 1$ to N_p :

Normalise weights: $\tilde{w}_{t_k}^i = w_{t_k}^i / \sum_{j=1}^{N_p} w_{t_k}^j$.

end for

if ESS < threshold:

Resample $\{s_{t_k}^i, \tilde{w}_{t_k}^i\}_{i=1}^{N_p}$ to obtain $\{s_{t_k}^i, \frac{1}{N_p}\}_{i=1}^{N_p}$.

end if

end for

How to approximate $p(s_{t'}|o_{t_{1:k}}; \theta)$ at an arbitrary time instance t' ?

- ▶ Assume $t' = t_k + \Delta t$.

$$\begin{aligned} p(s_{t_k+\Delta t}|o_{t_{1:k}}; \theta) &= \int p(s_{t_{1:k}}, s_{t_k+\Delta t}|o_{t_{1:k}}; \theta) ds_{t_{1:k}} \\ &= \int p(s_{t_{1:k}}|o_{t_{1:k}}; \theta) p(s_{t_k+\Delta t}|s_{t_k}; \theta) ds_{t_{1:k}}. \end{aligned}$$

- ▶ Draw samples $\{s_{t_k+\Delta t}^i\}_{i=1}^{N_p}$ from the dynamic process:

$$ds_t = f_\theta(s_t, t)dt + \sigma_\theta(s_t, t)dB_t.$$

- ▶ Approximate $p(s_{t_k+\Delta t}|o_{t_{1:k}}; \theta) \approx \frac{1}{N_p} \sum_{i=1}^{N_p} s_{t_k+\Delta t}^i$.

Two sets of experiments:

- ▶ Sequential state estimation (supervised training).
 - ▶ Benes-Daum filtering problem¹⁴.
 - ▶ Angular position prediction in a noisy physical pendulum model¹⁵.
- ▶ Observation prediction (unsupervised training).
 - ▶ Geometric Brownian motion⁴.
 - ▶ Stochastic Lorenz attractor³.

³Li et al., "Scalable gradients for stochastic differential equations", AISTATS, 2020.

⁴Deng et al., "Continuous Latent Process Flows", NeurIPS, 2021.

¹⁴Daum "Exact finite-dimensional nonlinear filters", IEEE Trans. Automatic Control, 1986.

¹⁵Sarkka, "Recursive Bayesian inference on stochastic differential equations", PhD thesis, Helsinki University of Technology, 2006.

Benes-Daum filtering problem:

$$ds_t = \tanh(s_t)dt + dB_t,$$
$$o_t \sim \mathcal{N}(s_t, \sigma^2).$$

- ▶ Estimate s_t at time t_k given observations $\{o_{t_0}, o_{t_1}, \dots, o_{t_k}\}$.
- ▶ Time stamps are sampled from a homogeneous Poisson process with intensity $\lambda \in \{0.5, 2\}$ in the interval $[0, 10]$.

Experiments: Benes-Daum filtering problem

Experimental results:

- ▶ Evaluated methods are compared by computing the RMSEs between their estimation of latent states and the ground-truth latent state.

Method	Poisson intensity	
	$\lambda=0.5$	$\lambda=2$
CD-DPF (ours)	1.33	0.689
Latent SDE ³	1.48	0.846
CLPF ⁴	1.51	0.879
CTPF ⁵	1.46	0.724

Filtering RMSE.

³Li et al., "Scalable gradients for stochastic differential equations", AISTATS, 2020.

⁴Deng et al., "Continuous Latent Process Flows", NeurIPS, 2021

⁵Deng et al., "Continuous-time Particle Filtering for Latent Stochastic Differential Equations", arXiv, 2209.00173.

Noisy pendulum system:

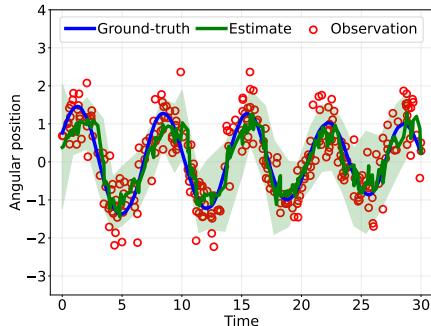
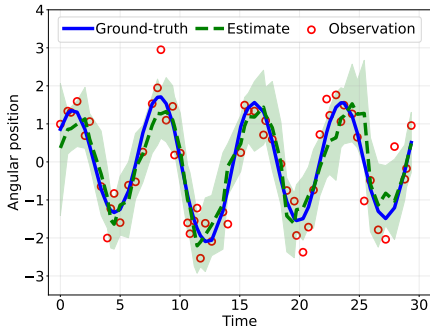
$$\begin{aligned}ds_t(1) &= s_t(2)dt, \\ds_t(2) &= -a^2 \sin(s_t(1))dt + b^{\frac{1}{2}}dB_t, \\o_t|s_t &\sim \mathcal{N}(s_t(1), \sigma^2).\end{aligned}$$

- ▶ Predict the angular position $s_{t_{k+1}}(1)$ of the pendulum at the next time stamp t_{k+1} given observations $\{o_{t_0}, o_{t_1}, \dots, o_{t_k}\}$.
- ▶ Time stamps are sampled from a homogeneous Poisson process with intensity $\lambda \in \{2.0, 10.0\}$ in the interval $[0, 30]$.

Experiments: angular position prediction

Experimental results:

- ▶ Shaded areas specify the 95% quantiles of empirical distributions given by the CD-DPF.



³Li et al., "Scalable gradients for stochastic differential equations", AISTATS, 2020.

⁴Deng et al., "Continuous Latent Process Flows", NeurIPS, 2021

⁵Deng et al., "Continuous-time Particle Filtering for Latent Stochastic Differential Equations", arXiv, 2209.00173.

Experiments: angular position prediction

Experimental results:

Method	Poisson intensity	
	$\lambda=2.0$	$\lambda=10.0$
CD-DPF (ours)	0.487	0.451
Latent SDE ³	0.749	0.612
CLPF ⁴	0.622	0.587
CTPF ⁵	0.501	0.489

Filtering RMSE.

³Li et al., "Scalable gradients for stochastic differential equations", AISTATS, 2020.

⁴Deng et al., "Continuous Latent Process Flows", NeurIPS, 2021

⁵Deng et al., "Continuous-time Particle Filtering for Latent Stochastic Differential Equations", arXiv, 2209.00173.

Experiments: observation prediction and interpolation

- ▶ Geometric Brownian motion

$$ds_t = \mu s_t dt + \sigma s_t dB_t, \\ s_{t_0} = s_0,$$

- ▶ Stochastic Lorenz Attractor

$$ds_t(1) = \eta \left(s_t(2) - s_t(1) \right) dt + \alpha(1) dB_t, \quad s_{t_0}(1) = x_0(1), \\ ds_t(2) = \left(s_t(1)(\rho - s_t(3)) - s_t(2) \right) dt + \alpha(2) dB_t, \quad s_{t_0}(2) = x_0(2), \\ ds_t(3) = \left(s_t(3)s_t(2) - \beta s_t(3) \right) dt + \alpha(3) dB_t, \quad s_{t_0}(3) = x_0(3).$$


Both experiments have Gaussian measurements:

$$o_t \sim \mathcal{N}(s_t, \sigma^2)$$

Experiments: observation prediction and interpolation

Geometric Brownian motion experimental results:

- ▶ Yellow curve: prediction. Black cross: observations. Purple and blue: training samples.




Experiments: observation prediction and interpolation


Stochastic Lorenz Attractor experimental results:

- ▶ We can discover the true system model from corrupted data.

Data



Samples



Some unresolved problems

- ▶ There are a number of unbiased estimators of the weights.
 - ▶ Unclear which one is better in what cases.

Some unresolved problems

- ▶ There are a number of unbiased estimators of the weights.
 - ▶ Unclear which one is better in what cases.
- ▶ Exact simulation is required.
 - ▶ No longer unbiased estimates if approximation or discretisation is needed in simulation.
 - ▶ Exact simulation techniques only apply to very limited and simple SDEs.

Some unresolved problems

- ▶ There are a number of unbiased estimators of the weights.
 - ▶ Unclear which one is better in what cases.
- ▶ Exact simulation is required.
 - ▶ No longer unbiased estimates if approximation or discretisation is needed in simulation.
 - ▶ Exact simulation techniques only apply to very limited and simple SDEs.
- ▶ System identification problem in unsupervised training.
 - ▶ Different latent dynamics and measurements can possibly interpret observations equally well but some may produce poor filtering results.

Some unresolved problems

- ▶ There are a number of unbiased estimators of the weights.
 - ▶ Unclear which one is better in what cases.
- ▶ Exact simulation is required.
 - ▶ No longer unbiased estimates if approximation or discretisation is needed in simulation.
 - ▶ Exact simulation techniques only apply to very limited and simple SDEs.
- ▶ System identification problem in unsupervised training.
 - ▶ Different latent dynamics and measurements can possibly interpret observations equally well but some may produce poor filtering results.

Thank you!


Appendix

Normalising Flows

Definition of normalising flows:

$$y = \mathcal{T}_\theta(x),$$

where \mathcal{T}_θ is required to be an invertible transformation.



Normalising Flows

Definition of normalising flows:

$$y = \mathcal{T}_\theta(x),$$

where \mathcal{T}_θ is required to be an invertible transformation.

Why invertible transformations?

- ▶ Invertibility allows density estimation (change of variable):


$$p(y) = p(x) \left| \det \frac{dy}{dx} \right|^{-1}$$

An Example of Normalising Flow: Coupling Layer¹⁶

Real-NVP

- ▶ Coupling layers.

$$x = [x_1, x_2] \quad y = [y_1, y_2]$$



¹⁶Ding et al. "Density Estimation Using Real NVP", ICLR, 2017.

An Example of Normalising Flow: Coupling Layer¹⁶

Real-NVP

- ▶ Coupling layers.

The special structure of coupling layers leads to triangular Jacobian matrix:


$$\begin{aligned} y_{1:d} &= x_{1:d} \\ y_{d+1:D} &= x_{d+1:D} \odot \exp(c(x_{1:d})) + t(x_{1:d}) \end{aligned}$$
$$\frac{\partial y}{\partial x} = \begin{bmatrix} \mathbb{I} & 0 \\ \frac{\partial y_{d+1:D}}{\partial x_{1:d}^T} & \text{diag}(\exp[c(x_{1:d})]) \end{bmatrix}$$

¹⁵Ding et al. "Density Estimation Using Real NVP", ICLR, 2017.

Conditional Coupling Layer

We use conditional coupling layer to construct conditional Real-NVP:

$$x = [x_1, x_2] \quad y = [y_1, y_2]$$




Standard coupling layer

Conditional Coupling Layer

We use conditional coupling layer to construct conditional Real-NVP:

$$x = [x_1, x_2] \quad y = [y_1, y_2]$$



Conditional coupling layer

Conditional Coupling Layer

- ▶ Conditional coupling layer:

$$y_{1:d} = x_{1:d}$$
$$y_{d+1:D} = x_{d+1:D} \odot \exp(c(x_{1:d}, o)) + t(x_{1:d}, o)$$

- ▶ Standard coupling layer:

$$y_{1:d} = x_{1:d}$$
$$y_{d+1:D} = x_{d+1:D} \odot \exp(c(x_{1:d})) + t(x_{1:d})$$

Conditional Coupling Layer

- ▶ Conditional coupling layer:

$$y_{1:d} = x_{1:d}$$
$$y_{d+1:D} = x_{d+1:D} \odot \exp(c(x_{1:d}, o)) + t(x_{1:d}, o)$$

Still invertible and lead to triangular Jacobian matrix:

$$\frac{\partial y}{\partial x} = \begin{bmatrix} \mathbb{I} & 0 \\ \frac{\partial y_{d+1:D}}{\partial x_{1:d}} & \text{diag}(\exp[c(x_{1:d}, o)]) \end{bmatrix}$$