Normalising Flow-based Differentiable Particle Filters

Xiongjie Chen and Yunpeng Li

University of Surrey, UK

Differentiable Particle Filters

Particle filters:

- Dynamic model $p(\mathbf{s}_t|\mathbf{s}_{t-1}, \mathbf{a}_t; \theta)$: transition of hidden state.
- Measurements model $p(\mathbf{o}_t|\mathbf{s}_t; \theta)$: likelihood of the observation.
- Track the marginal posterior $p(\mathbf{s}_t|\mathbf{o}_{1:t}, \mathbf{a}_{1:t}; \theta)$ or joint posterior $p(\mathbf{s}_{1:t}|\mathbf{o}_{1:t}, \mathbf{a}_{1:t}; \theta)$.

How can we jointly learn the parameter set θ and approximate the marginal posterior $p(\mathbf{s}_t|\mathbf{o}_{1:t}, \mathbf{a}_{1:t}; \theta)$ or joint posterior $p(\mathbf{s}_{1:t}|\mathbf{o}_{1:t}, \mathbf{a}_{1:t}; \theta)$?

Framework of Differentiable Particle Filters

By using deep learning techniques, differentiable particle filters (DPFs) provide a flexible way to learn the dynamic model and the measurement model.

- Build dynamic model and measurement model with neural networks [1, 2, 3]:
- Optimise the networks by minimising an objective function:
  - Supervised loss: RMSE, negative log likelihood [1, 2];
  - Semi-supervised loss: pseudo-likelihood [3].

Bottleneck of Current DPF frameworks

1. Only able to generate Gaussian prior [1];
2. Bootstrap particle filtering framework [2];
3. Do not admit valid probability densities in measurement models [1, 2].

Normalising Flows

Normalising flows are a family of invertible transformations, i.e. one-to-one mappings. In this work, we use the (conditional) Real-NVP [4].

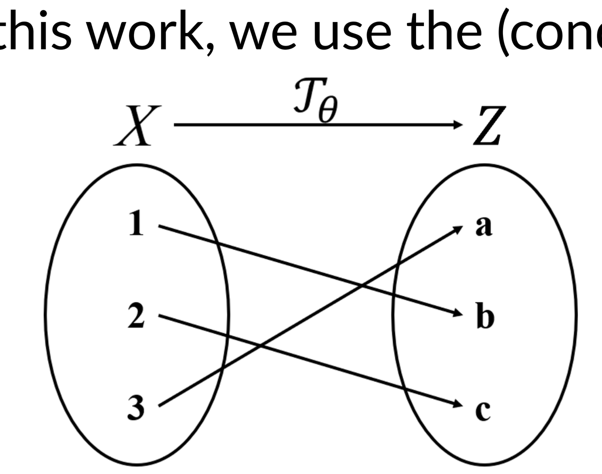

Why Normalising Flows?

- NFs can calculate exact probability densities;
  - Standard normalising flows $\mathcal{T}_\theta(\mathbf{x})$ can calculate the density $p(\mathbf{x})$;
  - Conditional normalising flows $\mathcal{G}_\theta(\mathbf{x}, \mathbf{y})$ can estimate the conditional density $p(\mathbf{x}|\mathbf{y})$.
- NFs can construct arbitrarily complex distributions.

Normalising Flows for Dynamic Models [5]

- With normalising flows, we can construct flexible dynamic models.

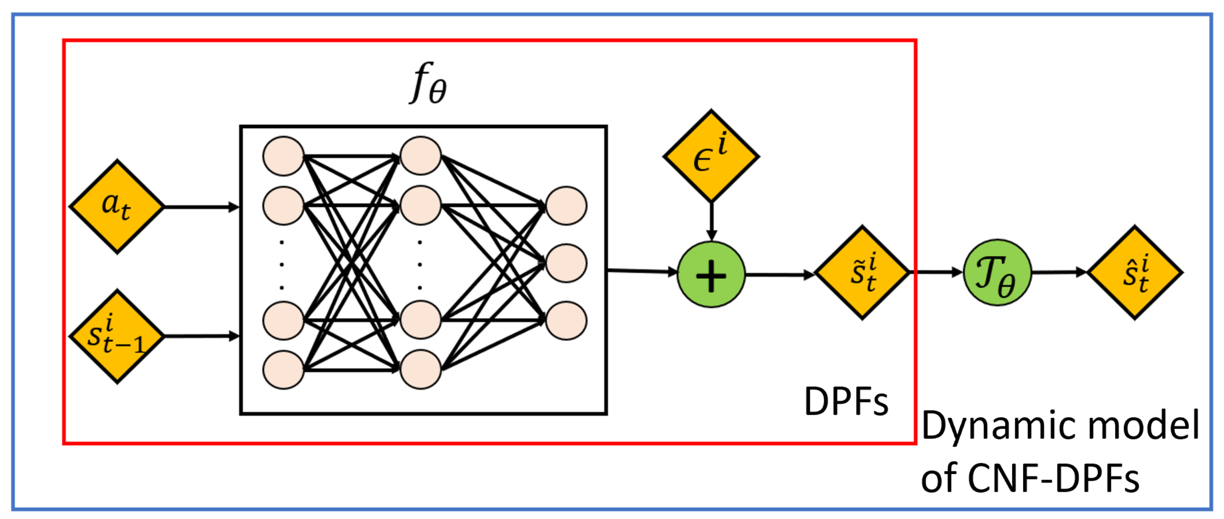

Proposals with Conditional Normalising Flows [5]

- Conditional normalising flows can move particles to areas closer to posterior by utilising information from observations.

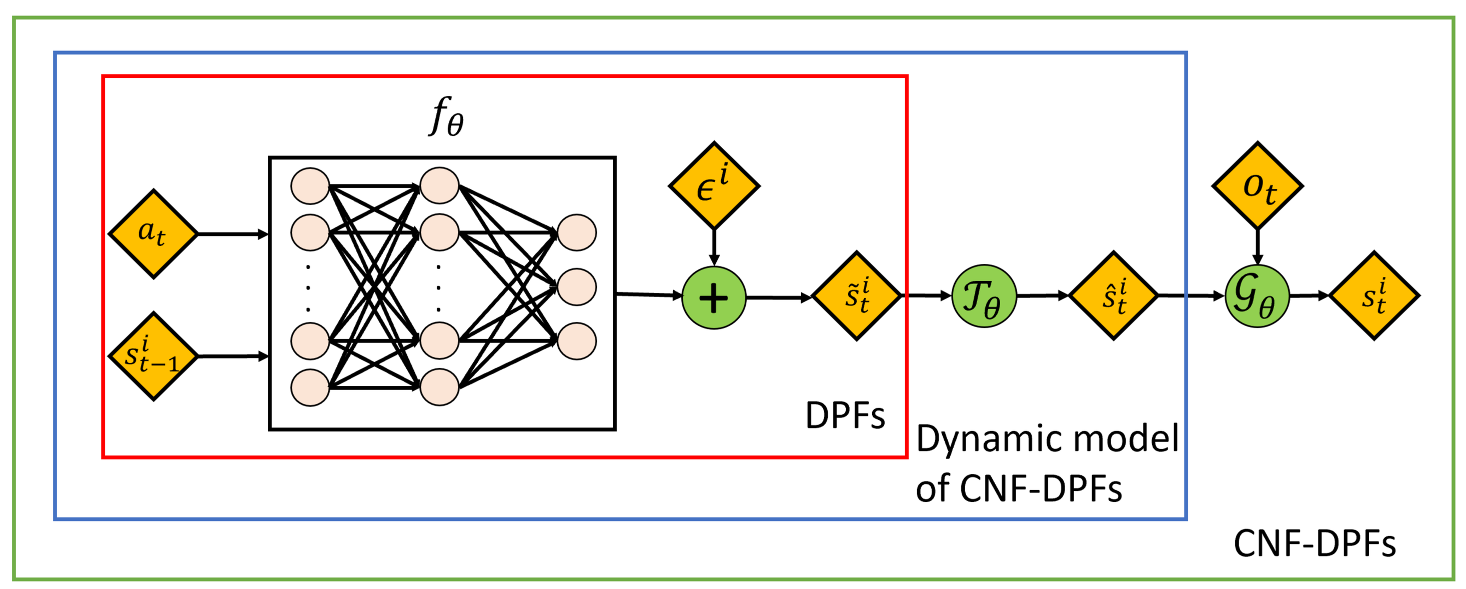

Conditional Normalising Flows-based Measurement Models [6]

- We can also construct measurement models using conditional normalising flows.

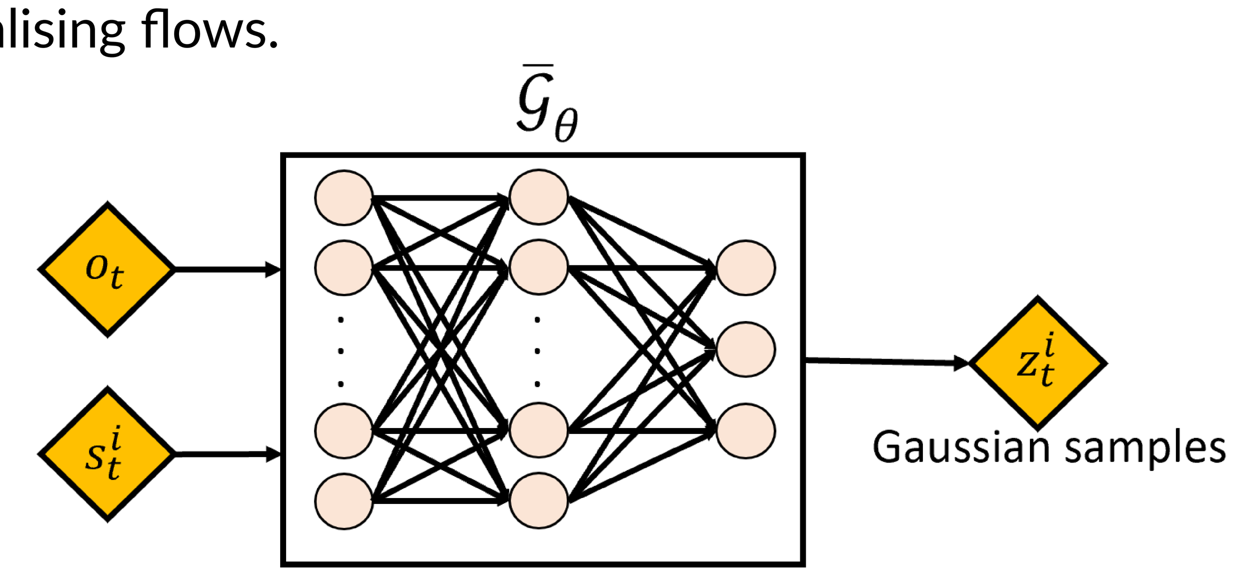

The likelihood of a measurement $\mathbf{o}_t$ given state $\mathbf{s}_t$ can be computed by using the change of variable formula:

$$p(\mathbf{o}_t|\mathbf{s}_t^i; \theta) = p_Z(\bar{\mathcal{G}}_\theta(\mathbf{o}_t, \mathbf{s}_t^i)) \left| \det \frac{\partial \bar{\mathcal{G}}_\theta(\mathbf{o}_t, \mathbf{s}_t^i)}{\partial \mathbf{o}_t} \right| ,$$

where the base distribution $p_Z(\cdot)$ of $\mathbf{z}_t^i$ can be user-specified and is often chosen as a simple distribution such as isotropic Gaussian.

Experiment Results

We evaluate the performance of our model in a visual tracking task, where the goal is to track the position of the red disk based on the observation images.

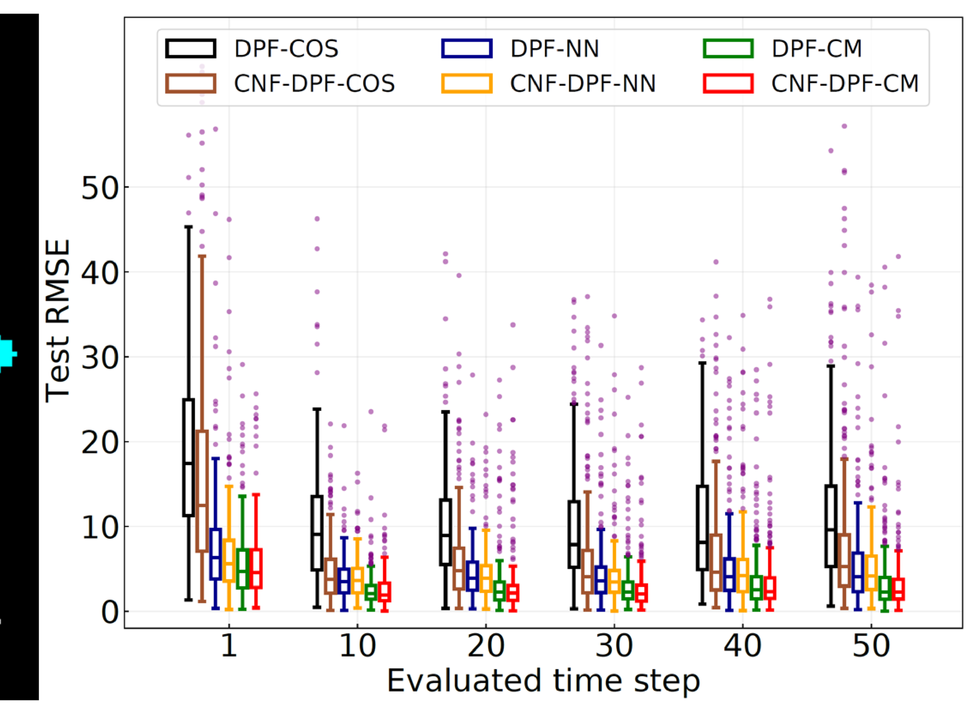

- The prefix "CNF-" indicates the dynamic model and proposal distribution are constructed with (conditional) normalising flows. The CNF-DPFs produced lower RMSEs compared with their counterparts;
- Different suffixes refer to different measurement models, the conditional normalising flow-based measurement model "-CM" exhibits the lowest RMSEs among all evaluated approaches.

Take-away Message

1. Normalising flows can be used to construct flexible dynamic models and proposal distributions;
2. The likelihood of measurements given states can be estimated by using conditional normalising flows;
3. Improved performances are observed through numerical experiments in a visual tracking task.

References

[1] P. Karkus, D. Hsu, and W. S. Lee. Particle filter networks with application to visual localization. In *Proc. Conf. Robot Learn. (CoRL)*, Zurich, Switzerland, 2018.

[2] R. Jonschkowski, D. Rastogi, and O. Brock. Differentiable particle filters: end-to-end learning with algorithmic priors. In *Proc. Robot.: Sci. and Syst. (RSS)*, Pittsburgh, Pennsylvania, 2018.

[3] H. Wen et al. End-to-end semi-supervised learning for differentiable particle filters. In *Proc. IEEE Int. Conf. Robot. Automat., (ICRA)*, Xi'an, China, 2021.

[4] C. Winkler et al. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019.

[5] X. Chen, H. Wen, and Y. Li. Differentiable particle filters through conditional normalizing flow. In *Proc. Int. Conf. on Inf. Fusion (FUSION)*, pages 1--6. IEEE, 2021.

[6] X. Chen and Y. Li. Conditional measurement density estimation in sequential Monte Carlo via normalizing flow. *arXiv preprint arXiv:2203.08617*, 2022.