

# Particle filters in high-dimensions

*Torben Sell*

May 23, 2018

1 Particle filtering

2 Transport maps

3 Future research

1 Particle filtering

2 Transport maps

3 Future research

Let  $\mathcal{X}$  be the state-space, and  $\mathcal{Y}$  be the observation space of a state space model. We assume the dynamics are as follows:

$$X_t = f(X_{t-1}) + \zeta_t$$

$$Y_t = h(X_t) + \eta_t,$$

where  $\zeta_t$  and  $\eta_t$  are independent random variables, and  $X_0 \sim \mu_0$ .

Let  $\mathcal{X}$  be the state-space, and  $\mathcal{Y}$  be the observation space of a state space model. We assume the dynamics are as follows:

$$\begin{aligned}X_t &= f(X_{t-1}) + \zeta_t \\ Y_t &= h(X_t) + \eta_t,\end{aligned}$$

where  $\zeta_t$  and  $\eta_t$  are independent random variables, and  $X_0 \sim \mu_0$ .

Given some observations  $y_t$ , we are interested in finding the filtering distribution  $\mu_t(x_t|y_{1:t})$ .

Let  $\mathcal{X}$  be the state-space, and  $\mathcal{Y}$  be the observation space of a state space model. We assume the dynamics are as follows:

$$\begin{aligned}X_t &= f(X_{t-1}) + \zeta_t \\ Y_t &= h(X_t) + \eta_t,\end{aligned}$$

where  $\zeta_t$  and  $\eta_t$  are independent random variables, and  $X_0 \sim \mu_0$ .

Given some observations  $y_t$ , we are interested in finding the filtering distribution  $\mu_t(x_t|y_{1:t})$ .

*Particle filtering* achieves this using a particle approximation to the true distribution.

# Particle filtering - applications

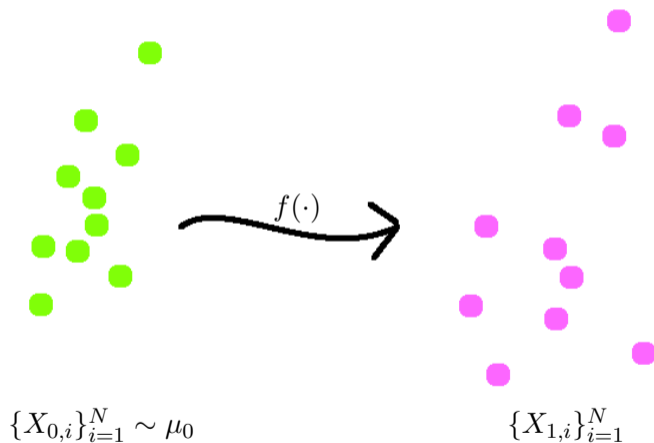
- Climate modelling
- Oceanography
- Multi-target tracking

# Particle filtering

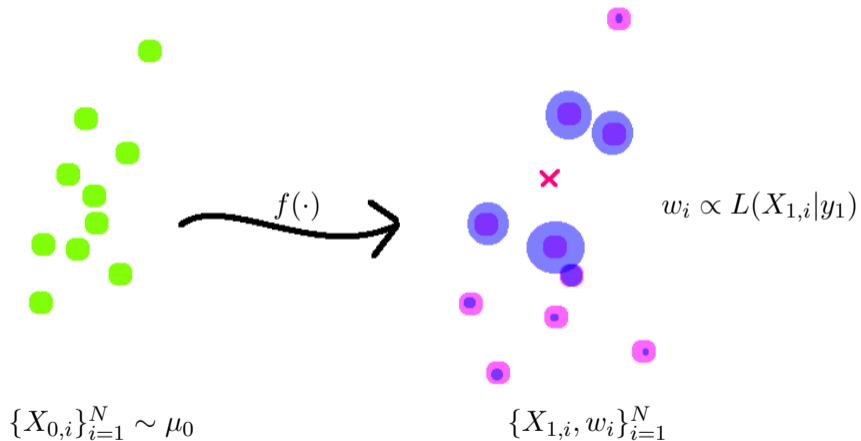


$$\{X_{0,i}\}_{i=1}^N \sim \mu_0$$

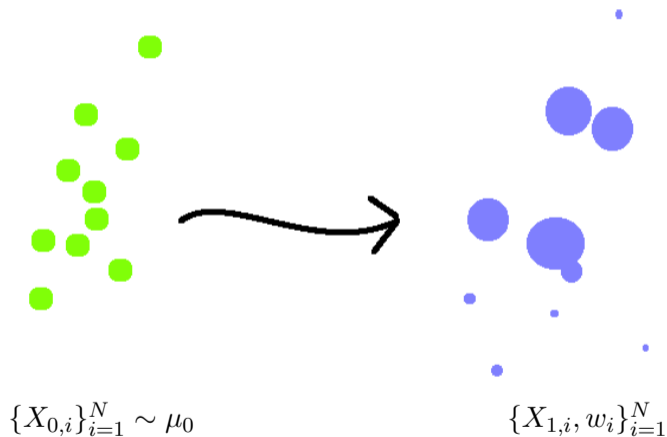
# Particle filtering



# Particle filtering



# Particle filtering



# Particle filtering - in high dimensions



$$\{X_{0,i}\}_{i=1}^N \sim \mu_0$$

$$\{X_{1,i}, w_i\}_{i=1}^N$$

# Particle filtering - the ideal case



$$\{X_{0,i}\}_{i=1}^N \sim \mu_0$$

$$\{X_{1,i}, w_i = \frac{1}{N}\}_{i=1}^N$$

1 Particle filtering

2 Transport maps

3 Future research

Ideally, one wants to find a transport map  $T_t : \mathcal{X} \rightarrow \mathcal{X}$  such that, if  $\{\hat{X}_i\}_{i=1}^N$  is a sample from the prior, then  $\{T_t(\hat{X}_i)\}_{i=1}^N$  is a sample from the posterior (the filtering distribution).

Ideally, one wants to find a transport map  $T_t : \mathcal{X} \rightarrow \mathcal{X}$  such that, if  $\{\hat{X}_i\}_{i=1}^N$  is a sample from the prior, then  $\{T_t(\hat{X}_i)\}_{i=1}^N$  is a sample from the posterior (the filtering distribution).

This may be achieved by finding a flow  $f$  which solves the PDE  $\partial_t \pi_t = -\nabla \cdot (\pi_t f)$ , where  $\pi_t$  is a curve of measures with  $\pi_0$  being the prior and  $\pi_1$  being the posterior. [Heng et al. 2015]

Then a transport map is given by the final location of a particle  $X_i$  under the flow  $f$ .

1 Particle filtering

2 Transport maps

3 Future research

To tackle the challenges which particle filters face in high-dimensional settings, promising research directions include:

- Improving on inaccurate (but fast) approximations obtained using Ensemble Kalman Filters.
- Exploit time-forgetting properties and space-forgetting properties to accelerate existing methods.
- ...

# Thank you for listening!

