Particle Flow Bayes' Rule

Xinshi Chen[1*], Hanjun Dai[1*], Le Song[1,2]
[1]Georgia Tech, [2]Ant Financial
(*equal contribution)

ICML 2019

Sequential Bayesian Inference

1. **Prior** distribution $\pi(x)$
2. **Likelihood** function $p(o|x)$
3. **Observations** $o_1, o_2, \dots, o_m$

   arrive sequentially

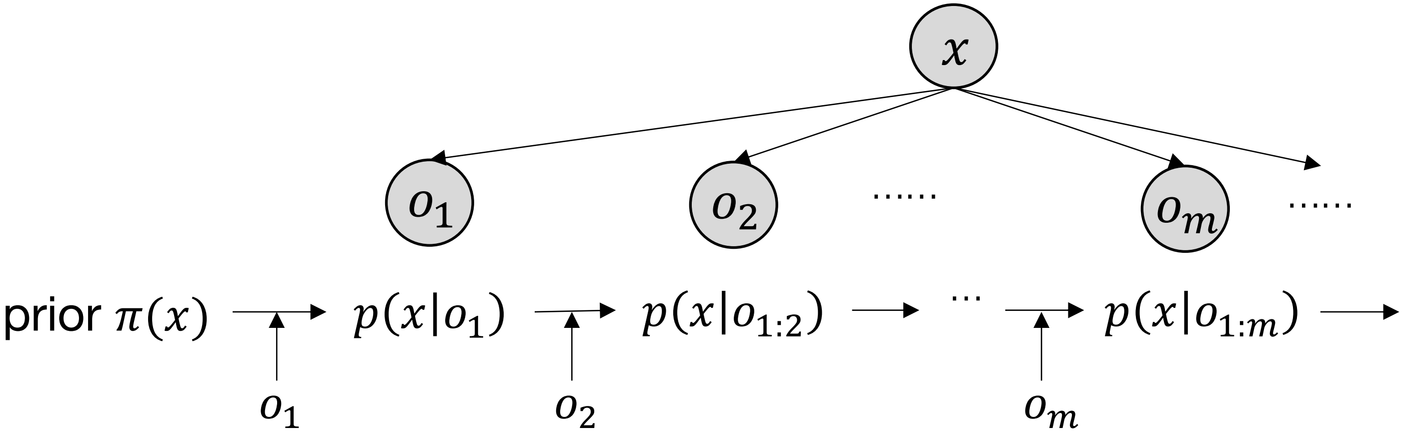

$$\underbrace{p(x|o_{1:m})}_{\text{updated posterior}} \propto \underbrace{p(x|o_{1:m-1})}_{\text{current posterior}} \underbrace{p(o_m|x)}_{\text{Likelihood Or prior}}$$

Need efficient online update!

Sequential Monte Carlo:

- N particles $\mathcal{X}_0 = \{x_0^1, \dots, x_0^N\}$ from prior $\pi(x)$
- Reweight the particles using likelihood
- Particle **degeneracy** problem

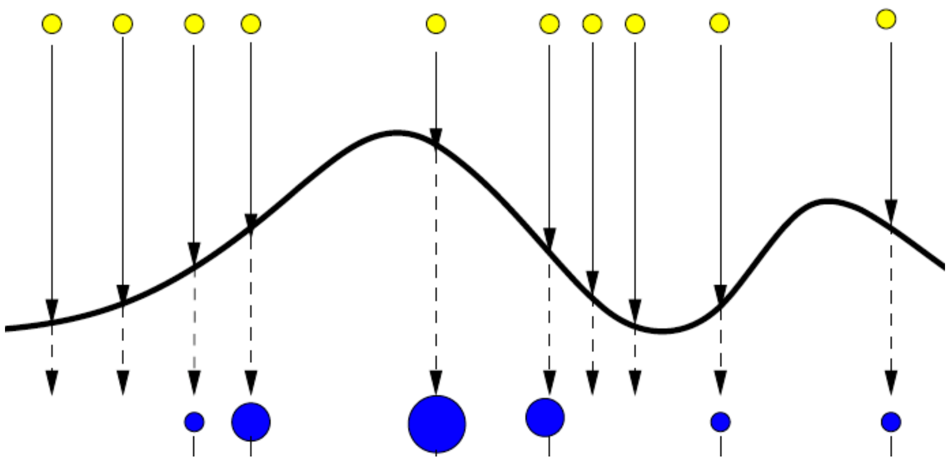

(Doucet et al., 2001)

Our Approach: Particle Flow

- N **particles** $\mathcal{X}_0 = \{x_0^1, \ldots, x_0^N\}$, **from prior** $\pi(x)$
- **Move particles through an ordinary differential equation (ODE)**

$$x(0) = x_0^n \quad \text{and} \quad \frac{dx}{dt} = f(\mathcal{X}_0, o_1, x(t), t)$$

$$\Longrightarrow \quad \text{solution } x_1^n = x(T)$$

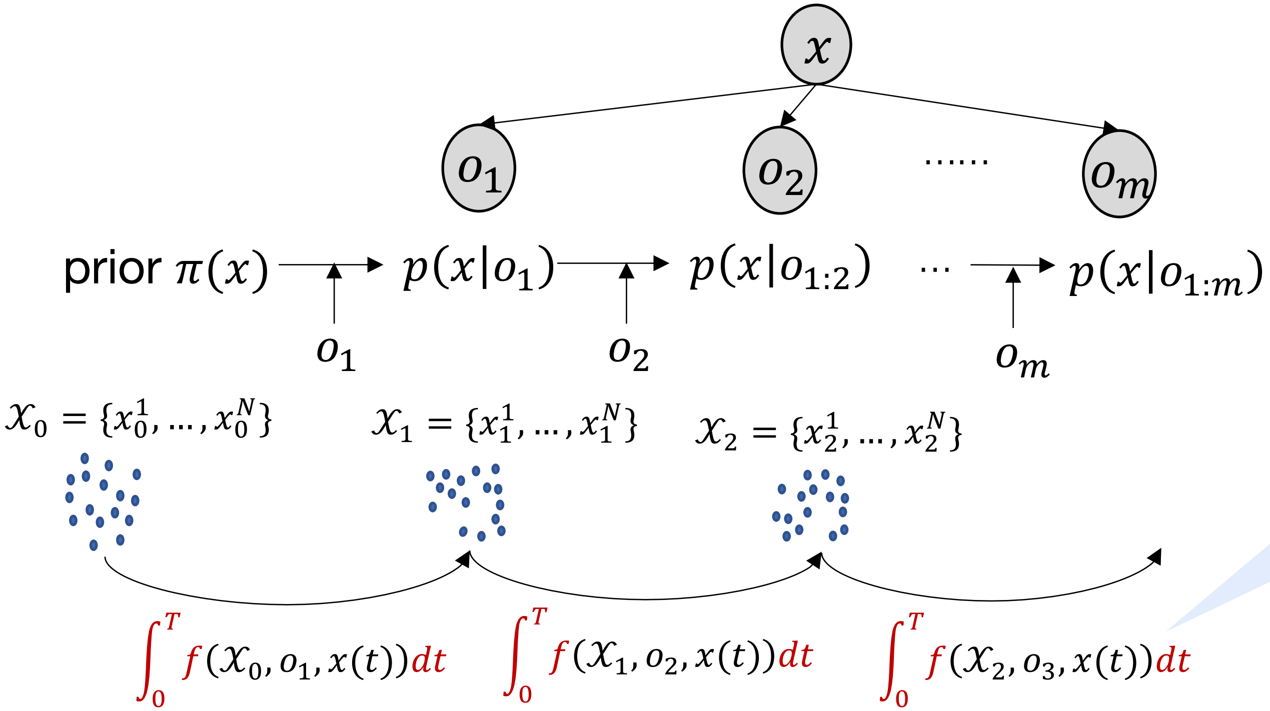

Does a **unified** flow velocity f exist?
Does **P**article **F**low **B**ayes' **R**ule **(PFBR)** exist?

Our Approach: Particle Flow

- **Move particles to next posterior through an ordinary differential equation (ODE)**

$$x(0) = x_0^n \quad \text{and} \quad \frac{dx}{dt} = f(\mathcal{X}_0, o_1, x(t), t)$$

$$\Longrightarrow \quad \text{solution } x_1^n = x(T)$$

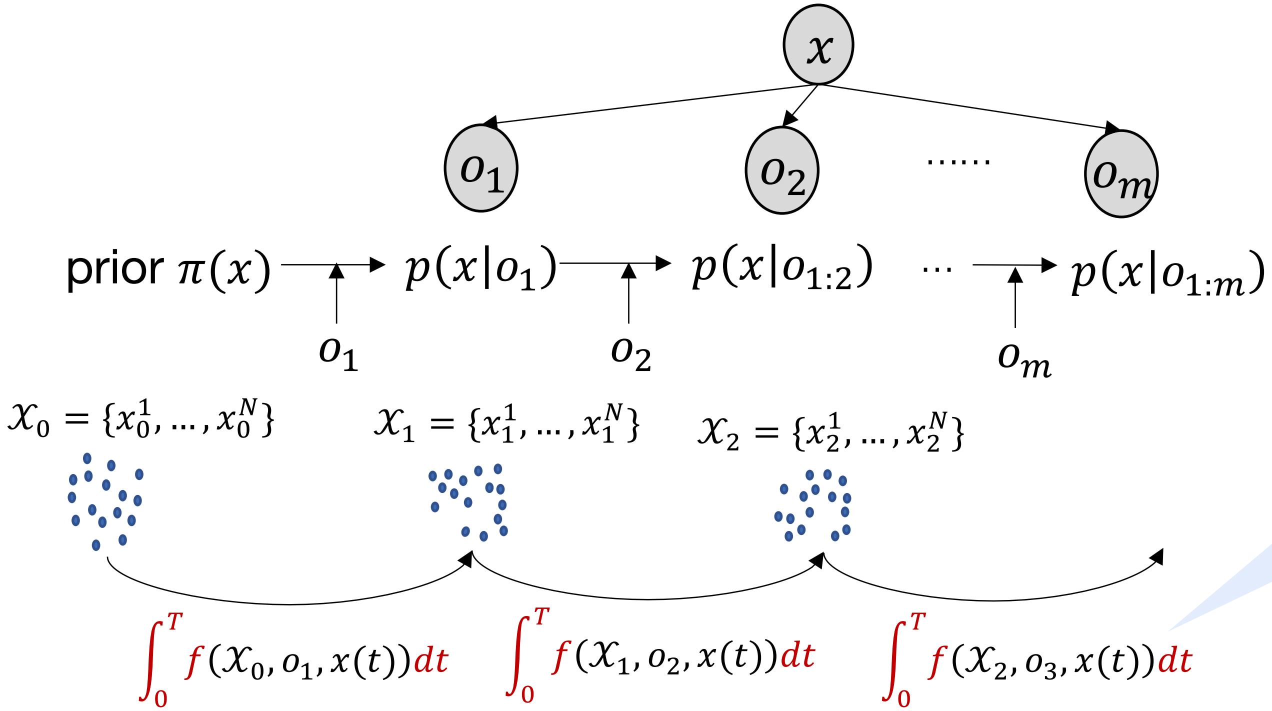

Does a **unified** flow velocity f exist?
Does **P**article **F**low **B**ayes' **R**ule **(PFBR)** exist?

Yes!!!

$$f := \nabla_x \log \pi(x) p(o|x) - w^*(\pi(x), p(o|x), x, t)$$

Existence of Particle Flow Bayes' Rule

Langevin dynamics

$$dx(t) = \nabla_x \log \pi(x) p(o|x)\, dt + \sqrt{2}\, \mathbf{dw(t)}$$

- ✓ density $q(x,t)$ converges to posterior $p(x|o)$
- ✗ stochastic flow

Fokker-Planck Equation + Continuity Equation

deterministic, closed-loop

$$dx(t) = \nabla_x \log \pi(x) p(o|x) - \nabla_x \log \mathbf{q(x,t)}\, dt$$

- ✓ density $q(x,t)$ converges to posterior $p(x|o)$
- ✓ deterministic flow
- ✗ closed-loop flow: depends on $q(x,t)$

Optimal control theory: closed-loop to open loop

deterministic, open-loop

$$dx(t) = \nabla_x \log \pi(x) p(o|x) - w^*(\pi(x), p(o|x), x, t) dt$$

- ✓ density $q(x,t)$ converges to posterior $p(x|o)$
- ✓ deterministic flow
- ✓ open-loop flow

Parameterization

The unified flow velocity is in form of:

$$\boldsymbol{f}(\boldsymbol{\pi}(\boldsymbol{x}), \boldsymbol{p}(\boldsymbol{o}|\boldsymbol{x}), \boldsymbol{x}, \boldsymbol{t}) := \nabla_x \log \pi(x) p(o|x) - w^*(\pi(x), p(o|x), x, t)$$

$$\frac{\boldsymbol{dx}}{\boldsymbol{dt}} = \boldsymbol{f}(\mathcal{X}, \boldsymbol{o}, \boldsymbol{x}(\boldsymbol{t}), \boldsymbol{t}) := \boldsymbol{h}\left(\frac{\mathbf{1}}{\boldsymbol{N}} \sum_{\boldsymbol{n}=1}^{\boldsymbol{N}} \boldsymbol{\phi}(\boldsymbol{x}^{\boldsymbol{n}}), \boldsymbol{o}, \boldsymbol{x}(\boldsymbol{t}), \boldsymbol{t}\right)$$

Deep set

h **neural networks**

Experiment 1: Multimodal Posterior

Gaussian Mixture Model

- prior $x_1, x_2 \sim \mathcal{N}(0,1)$
- observations $\mathrm{o}|x_1, x_2 \sim \frac{1}{2}\mathcal{N}(x_1, 1) + \frac{1}{2}\mathcal{N}(x_1 + x_2, 1)$
- With $(x_1, x_2) = (1, -2)$, **the resulting posterior $\boldsymbol{p(x|o_1, \dots, o_m)}$ will have two modes:**

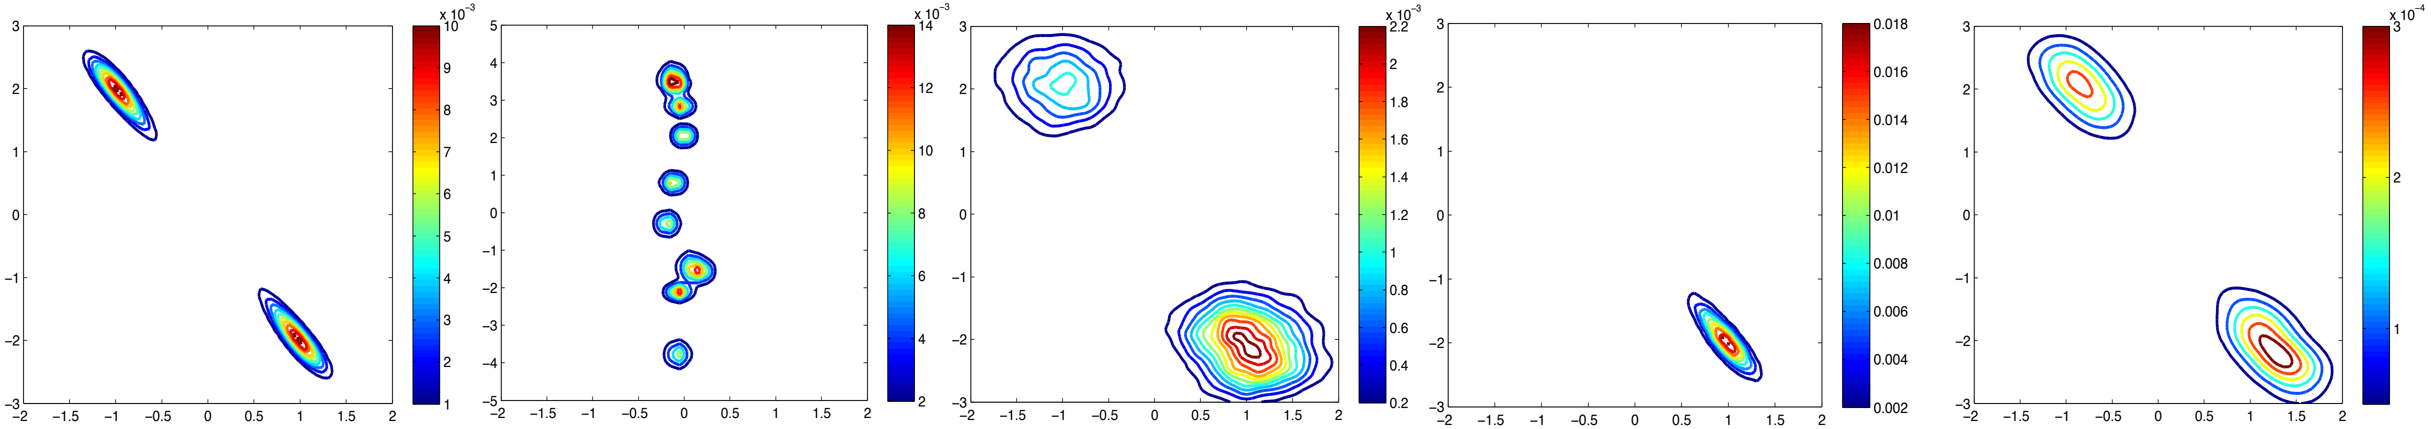

(a) True posterior (b) Stochastic Variational Inference (c) Stochastic Gradient Langevin Dynamics (d) Gibbs Sampling (e) One-pass SMC

Experiment 1: Multimodal Posterior

PFBR vs one-pass SMC

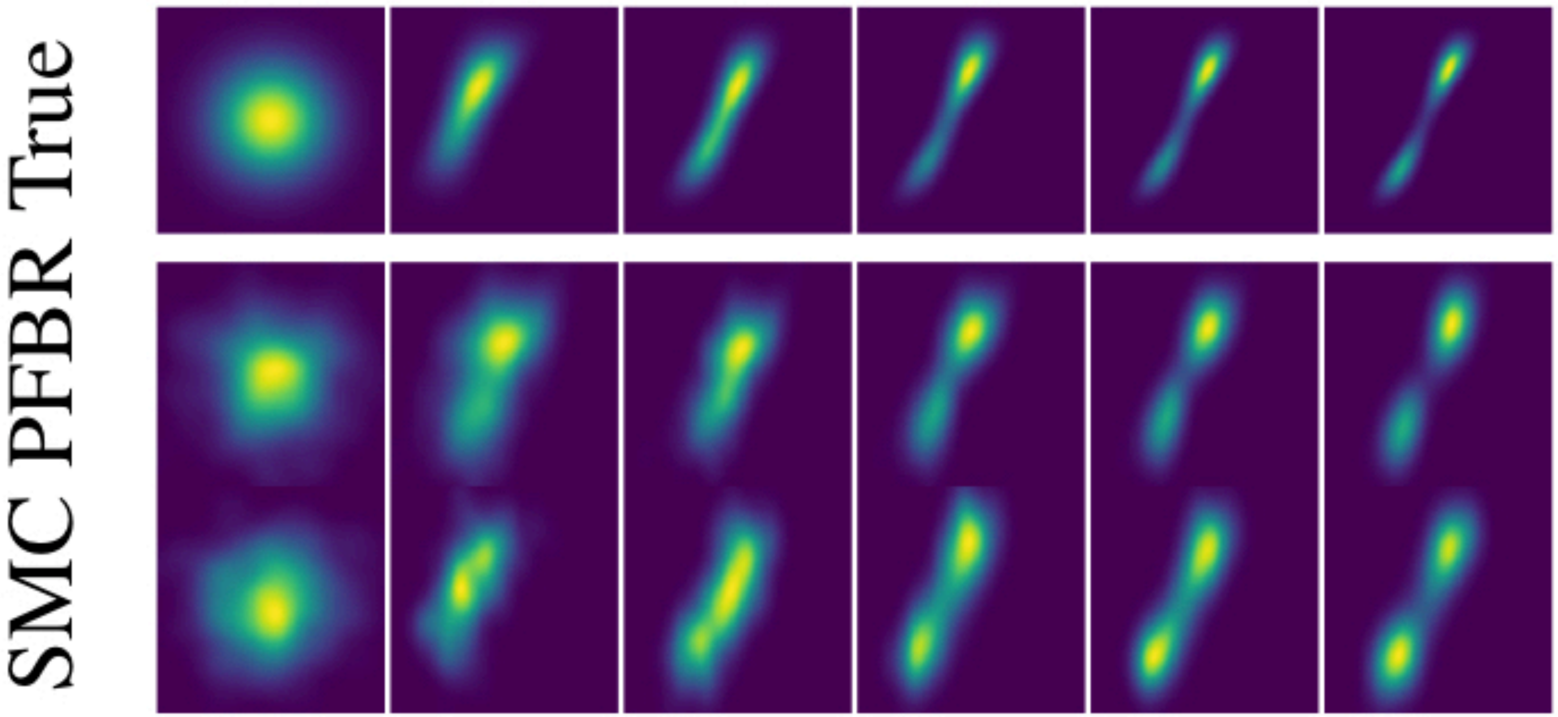

Visualization of the evolution of posterior density from left to right.

Experiment 2: Efficiency in #Particles

	Algo	#particles	cpu time (s)	gpu time (s)	cross-entropy
Our Approach	PFBR	256	0.23	0.26	16.56
	SMC	256	0.07	0.02	26.78
	ASMC-mlp	256	0.17	0.07	19.66
	ASMC-gru	256	0.18	0.07	19.38
	ASMC-mlp	4096	2.23	0.25	17.63
	ASMC-gru	4096	2.26	0.26	17.24
	SMC	8192	3.87	0.12	17.60

Comparison to SMC and ASMC (Autoencoding SMC, Filtering Variational Objectives, and Variational SMC) (Le et al., 2018; Maddison et al., 2017; Naesseth et al., 2018).

Thanks!

Poster: Pacific Ballroom #218, Tue, 06:30 PM

Contact: xinshi.chen@gatech.edu