Convergence of particle filters and relation to data assimilation

Dan Crisan

Imperial College London

Discussion Meeting on Non Linear Filtering and Data Assimilation
International Centre for Theoretical Sciences
Tata Institute of Fundamental Research
8-12 January 2014

Mini-course synopsis

- Day I
  - Stochastic Filtering (SF) in discrete time/Probabilistic formulation of data assimilation (DA)
  - Algorithms: a common language for SF and DA
  - Particle Filters
- Day II
  - The Framework in detail
  - The Recurrence formula for π_t
  - Convergence of approximations to π_t
- Day II
  - Particle Filters
  - Why are the high-dimensional problem hard ?
  - Application to the observed Navier-Stokes equation
  - Future work

$(X, Y) = \{(X_t, Y_t), t \geq 0\}$

- X the signal process - "hidden component"
- Y the observation process - "the data" - $Y_t = f(X_t, \text{"}noise\text{"})$.

The Stochastic Filtering/Data Assimilation problem: Find the conditional distribution of the *signal* X_t given $\mathcal{Y}_t = \sigma(Y_s, s \in [0, t])$, i.e.,

$$\pi_t(A) = \mathbb{P}(X_t \in A | \mathcal{Y}_t), \qquad t \geq 0.$$

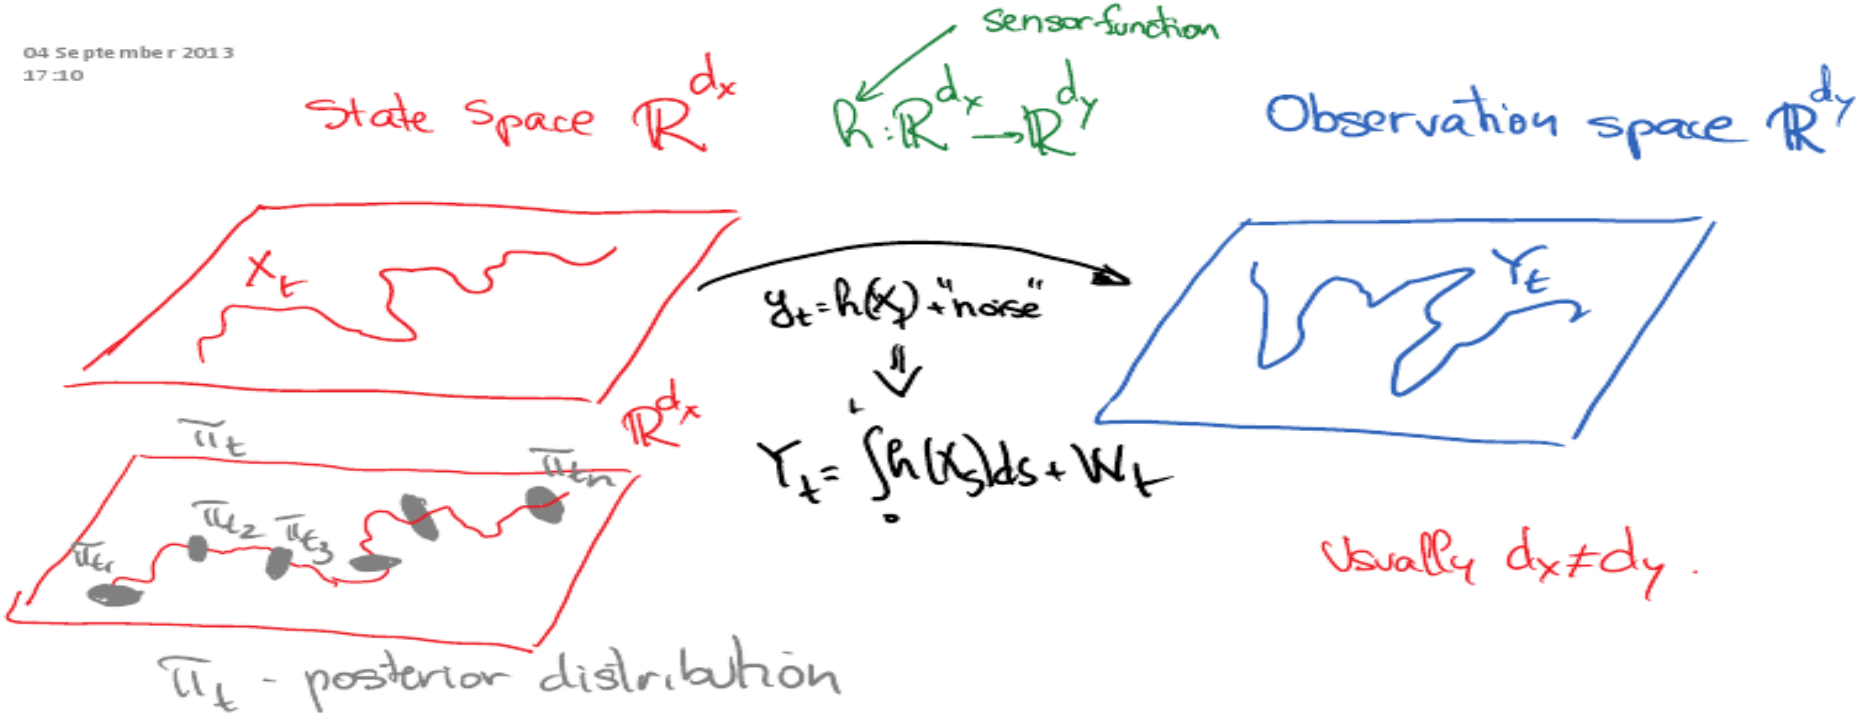

The filtering problem in discrete time
Probabilistic formulation of DA

The signal process: $X = \{X_t; t \in \mathbb{N}\}$, Markov chain with state space $\mathbb{R}^{d_x}$,

$$X_0 \sim \pi_0(dx_0), \quad \mathbb{P}(X_t \in A | X_{t-1} = x_{t-1}) = K_t(x_{t-1}, A).$$

Example ($d_x = 1$):

$$X_t = b(X_{t-1}) + B_t, \qquad B_t \sim N(0,1) \text{ i.i.d.}$$

$$K(x_{t-1}, A) = \int_A \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x_t - b(x_{t-1}))^2}{2}\right) dx_t$$

The observation process: Y associated stochastic process with state space $\mathbb{R}^{d_y}$ such that

$$\Pr\left(Y_t \in B | \mathcal{F}_t^X\right) = \Pr(Y_t \in B | X_t = x_t) = \int_B g_t(y_t, x_t)\, dy_t.$$

Example ($d_y = 1$):

$$Y_t = h(X_t) + V_t, \qquad V_t \sim N(0,1) \text{ i.i.d.}$$

$$g_t(y_t, x_t) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(y_t - h(x_t))^2}{2}\right)$$

The filtering problem consists in computing:

$$\pi_t \text{ - the conditional distribution of } X_t \text{ given } \left\{ Y_{[0,t]} = y_{[0,t]} \right\}$$

where $Y_{[0,t]} \triangleq (Y_0, ..., Y_t)$, $y_{[0,t]} \triangleq (y_0, ..., y_t)$.

Bayes' recursion.

$$\begin{array}{ll} \text{Prediction Step} & p_t = \pi_{t-1} K_t \\ \text{Updating Step} & \frac{d\pi_t}{dp_t} = C_t^{-1} g_t \end{array},$$

where

- $C_t \triangleq \int_{\mathbb{R}^{d_X}} g_t(y_t, x)\, p_t(dx)$.
- If μ is a measure and K is a kernel, then

$$\mu K(A) \triangleq \int \mu(dx)\, K(x, A).$$

Data assimilation (DA) is the process by which observations are incorporated into a computer model of a real system. Applications of data assimilation arise in many fields of geosciences, perhaps most importantly in weather forecasting and hydrology.

DA proceeds by alternating between forecast and analysis cycles:

- In each analysis cycle, observations of the current (and possibly past) state of a system are combined with the results from a prediction model (the forecast) to produce an analysis. The analysis step is typically performed either in form of a 'best estimate' or in terms of approximating conditional distributions.
- The model is then advanced in time and its result becomes the forecast in the next analysis cycle.

- Both DA and stochastic filtering are dealing with the same problem of merging models with partial observations
- DA has stronger focus on algorithms for large scale problems and large data sets
- Stochastic filtering stronger focus on asymptotic behavior and consistency
- Both fields are moving towards each other
- Situation is similar to that of machine learning versus statistics

Particle Filters/Sequential Monte Carlo methods

- Algorithms to approximate π_t using discrete measures of the form [1]

$$\sum_i \alpha_i \delta_{V_i},$$

  i.e., empirical distributions associated with a set of (random) particles with masses $\alpha_1, \alpha_2, ...$, and positions $V_1, V_2, ...$, respectively, in the state space of X.
- Recursive algorithms: The approximation for π_t and Y_{t+1} are the only information used in order to obtain the approximation for π_{t+1}. In other words, the information gained from $Y_1, ..., Y_t$ is embedded in the current approximation.
- Quite often:

$$\pi_t^n = \frac{1}{n} \sum_{k=1}^{n_t} \delta_{V_k^t}$$

  Consequently $E[n_t] = n$, with many of the existing algorithms keeping the number of particles constant: $n_t \equiv n$.

[1] δ_x is the Dirac delta distribution concentrated at x: $\delta_x(A) = \mathbf{1}_A(x)$.

- The generic SMC method involves sampling from the prior distribution of the signal and then using a weighted bootstrap technique (or equivalent) with weights defined by the likelihood of the most recent observation data.

- N. J. Gordon, D. J. Salmond, and A. F. M. Smith, *Novel approach to nonlinear/non-gaussian bayesian state estimation*, IEE Proceedings on Radar and Signal Processing, vol. 140, pp. 107–113, 1993.
- A. Doucet, N. de Freitas, N. Gordon, *Sequential Monte Carlo methods in practice*, Stat. Eng. Inf. Sci., Springer, New York, 2001.
- A. Bain, D. Crisan, Chapter 10, *Fundamentals of stochastic filtering*. Stochastic Modelling and Applied Probability, 60. Springer, New York, 2009.
- D. Crisan, B. Rozovskii, *The Oxford handbook of nonlinear filtering*, Oxford Univ. Press, Oxford, 2011.
- D. Crisan and S. Ortiz-Latorre, *A Kusuoka-Lyons-Victoir particle filter*, 2013, Proceedings of the Royal Society A, 2013.
- A. Beskos, D. Crisan, A Jasra., N. Whiteley, *Error Bounds and Normalizing Constants for Sequential Monte Carlo in High Dimensions*, 2013.
- A. Beskos, D. Crisan, A Jasra., *On the Stability of Sequential Monte Carlo Methods in High Dimensions*, 2013.

A common language

Ensemble-based algorithms in both areas fit into the following framework:

particle approximations	Gaussian approximations
$\pi_t^N \xrightarrow[\text{model}]{\text{mutation}} p_{t+\delta}^N \xrightarrow[\{Y_s\}_{s\in[t,t+\delta]}]{\text{selection}} \pi_{t+\delta}^N$	$\pi_t^N \xrightarrow[\text{model}]{\text{forecast}} p_{t+\delta}^N \xrightarrow[\{Y_s\}_{s\in[t,t+\delta]}]{\text{assimilation}} \pi_{t+\delta}^N$

The approximations *appear to be* different:

particle approximations	Gaussian approximations
$(\underbrace{a_j(t)}_{\text{weight}}, \underbrace{v_j^1(t), \ldots, v_j^d(t)}_{\text{position}})_{j=1}^N$	$(\underbrace{a_j(t)}_{\text{weight}}, \underbrace{v_j^1(t), \ldots, v_j^d(t)}_{\text{mean}}, \underbrace{\omega_j^{11}(t), \ldots, \omega_j^{dd}(t)}_{\text{covariance matrix}})_{j=1}^N$
$\pi_t \rightsquigarrow \pi_t^N = \sum_{j=1}^N a_j(t)\, \delta_{v_j(t)}$	$\pi_t \rightsquigarrow \pi_t^N = \sum_{j=1}^N a_j(t)\, N\left(v_j(t), \omega_j(t)\right)$

but the stored information can modelled by N stochastic processes

$$\{p_i(t), t > 0\} \quad i = 1, ..., N, \quad p_i(t) \in \mathbb{R}^M.$$

- We think of the processes p_i as the trajectories of N (generalized) particles/ensemble members.
- Typically $M > d_x$, where d_x is the dimension of the state space.

$$\pi_t^N = \Lambda_t^N(p_i(t), t > 0 \;\; i = 1, ..., N).$$

- Generalized particle filters:
  - classical particle filters
  - Gaussian approximations
  - wavelets
  - grid methods
- The measure of the approximating error is important:

$$\sup_{\{\varphi \in C_b, \;\; \|\varphi\| \leq 1\}} \mathbb{E}\left[\left|\pi_t^N(\varphi) - \pi_t(\varphi)\right|\right], \quad \hat{\pi}_t - \hat{\pi}_t^N, \quad \|\pi_t^N - \pi_t\|_{TV}.$$

Consider the 1-dimensional Benes filter:

$$\begin{aligned} dX_t &= \mu\sigma \tanh\left(\frac{\mu X_t}{\sigma}\right) dt + \sigma dV_t \\ dY_t &= (h_1 X_t + h_2)dt + dU_t, \end{aligned}$$

Then

$$\rho_t \simeq w^+ \mathcal{N}(A_t^+/(2B_t), 1/(2B_t)) + w^- \mathcal{N}(A_t^-/(2B_t), 1/(2B_t)),$$

where

$$\begin{aligned} w_t^{\pm} &\triangleq \exp\left((A_t^{\pm})^2/(4B_t)\right) / (\exp\left((A_t^+)^2/(4B_t)\right) + \exp\left((A_t^-)^2/(4B_t)\right)) \\ A_t^{\pm} &\triangleq \pm\frac{\mu}{\sigma} + h_1\Psi_t + \frac{h_2 + h_1 x_0}{\sigma \sinh(h_1\sigma t)} - \frac{h_2}{\sigma}\coth(h_1\sigma t), \\ B_t &\triangleq \frac{h_1}{2\sigma}\coth(h_1\sigma t), \\ \Psi_t &\triangleq \int_0^t \frac{\sinh(h_1\sigma s)}{\sinh(h_1\sigma t)} dW_s, \end{aligned}$$

The framework in details

Let the signal $X = \{X_t,\ t \in \mathbb{N}\}$ be a stochastic process with values in $\mathbb{R}^{d_x}$. Let $\mathcal{F}_t^X$ be the σ-algebra generated by the process, i.e.,

$$\mathcal{F}_t^X \triangleq \sigma(X_s,\ s \in [0,t]).$$

We assume that X is a Markov chain. That is, for all $t \in \mathbb{N}$ and $A \in \mathcal{B}(\mathbb{R}^{d_x})$,

$$\mathbb{P}\left(X_t \in A \mid \mathcal{F}_{t-1}^X\right) = \mathbb{P}\left(X_t \in A \mid X_{t-1}\right). \tag{1}$$

The transition kernel of the Markov chain X is the function $K_t(\cdot,\cdot)$ defined for all $t \in \mathbb{N}$ and $x \in \mathbb{R}^{d_x}$,

$$K_t(x, A) = \mathbb{P}(X_t \in A \mid X_{t-1} = x). \tag{2}$$

The transition kernel K_t is required to have the following properties:

- $K_t(x,\cdot)$ is a probability measure for all $t \in \mathbb{N}$ and $x \in \mathbb{R}^{d_x}$.
- $K_t(\cdot, A)$ is a measurable function for all $t \in \mathbb{N}$ and A Borel set.

The distribution of X is uniquely determined by its initial distribution and its transition kernel. Let us denote by q_t the distribution of the random variable X_t,

$$q_t(A) \triangleq \mathbb{P}(X_t \in A).$$

Then, from (2), it follows that q_t satisfies the recurrence formula

$$q_t = K_t q_{t-1}, \ t \geq 0,$$

Hence, by induction it follows that $q_t = K_{t-1} \dots K_1 K_0 q_0, \ t > 0$.

Example ($d_x = 1$):

$$X_t = b(X_{t-1}) + B_t, \qquad B_t \sim N(0,1) \text{ i.i.d.}$$

$$P(X_t \in dx_t, ..., X_0 \in dx_0) = \prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x_i - b(x_{i-1}))^2}{2}\right) dx_i \pi_0(dx_0)$$

$$\begin{aligned} q_t(A) &= \int_A \int_{\mathbb{R}} \int_{\mathbb{R}} \dots \int_{\mathbb{R}} \prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x_i - b(x_{i-1}))^2}{2}\right) dx_i \pi_0(dx_0) \\ &= \int_A \int_{\mathbb{R}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x_t - b(x_{t-1}))^2}{2}\right) q_{t-1}(dx_{t-1}) dx_t \end{aligned}$$

Exercise 1. For arbitrary $\varphi \in B(\mathbb{R}^d)$ and $t \geq 0$, define $K_t\varphi$ as

$$K_t\varphi(x) = \int_{\mathbb{R}^d} \varphi(y) K_t(x, dy).$$

i. Prove that $K_t\varphi$ is a measurable function for any $t \geq 0$.

ii. Prove that $K_t q_{t-1}$ is a probability measure for any $t \geq 0$.

iii. Prove that, for any $\varphi \in B(\mathbb{R}^d)$ and $t > 0$, we have

$$K_t q_{t-1}(\varphi) = q_{t-1}(K_t\varphi),$$

hence in general

$$q_t(\varphi) = q_0(\varphi_t), \ t > 0,$$

where $\varphi_t = K_0 K_1 \ldots K_{t-1}\varphi$.

Let the observation process $Y = \{Y_t,\ t \in \mathbb{N}\}$ be an $\mathbb{R}^{d_y}$-valued stochastic process defined as follows

$$Y_t \triangleq h(X_t) + W_t,\ t > 0, \tag{3}$$

and $Y_0 = 0$. In (3), $h : \mathbb{R}^{d_x} \to \mathbb{R}^{d_y}$ is a Borel-measurable function and for all $t \in \mathbb{N}$, W_t are mutually independent random vectors with laws absolutely continuous with respect to the Lebesgue measure λ on $\mathbb{R}^m$. We denote by $g(t, \cdot)$ the density of W_t with respect to λ and we further assume that $g(t, \cdot) \in B(\mathbb{R}^d)$ and is a strictly positive function.

Example ($d_x = d_y = 1$):

$$X_t = b(X_{t-1}) + B_t, \qquad B_t \sim N(0,1) \text{ i.i.d.}$$
$$Y_t = h(X_t) + V_t, \qquad V_t \sim N(0,1) \text{ i.i.d.}$$

$$P(Y_t \in dy_t, ..., Y_0 \in dy_1, X_t \in dx_t, ..., X_0 \in dx_0)$$
$$= \prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(y_i - h(x_{i-1}))^2}{2}\right) dy_i$$
$$\times \prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(x_i - b(x_{i-1}))^2}{2}\right) dx_i \pi_0(dx_0)$$

The filtering problem consists of computing the conditional distribution of the signal given the observation data from time 0 up to the current time i.e. computing the probability measure π_t, where

$$\pi_t(A) \triangleq \mathbb{P}(X_t \in A \mid Y_{0:t} = y_{0:t}), \tag{4}$$
$$\pi_t f = \mathbb{E}\left[f(X_t) \mid Y_{0:t} = y_{0:t}\right],$$

where $Y_{0:t}$ is the random vector $Y_{0:t} \triangleq (Y_0, Y_1, \ldots, Y_t)$[2] and $y_{0:t} \triangleq (y_0, y_1, \ldots, y_t) \in (\mathbb{R}^m)^{t+1}$.
We also introduce p_t, $t > 0$ the *predicted* conditional probability measures defined by

$$p_t(A) \triangleq \mathbb{P}\left(X_t \in A \mid Y_{0:t-1} = y_{0:t-1}\right),$$
$$p_t f = \mathbb{E}\left[f(X_t) \mid Y_{0:t-1} = y_{0:t-1}\right].$$

In the statistics and engineering literature the probability q_t is commonly called the *prior distribution* of the signal X_t, whilst π_t is called the (Bayesian) *posterior distribution*.

[2] $\{Y_{0:t}, t \in \mathbb{N}\}$ is the path process associated to the observation process $Y = \{Y_t,\ t \in \mathbb{N}\}$. That is, $\{Y_{0:t}, t \in \mathbb{N}\}$ records the entire history of Y up to time t, not just its current value.

The Recurrence Formula for π_t

The following lemma gives the density of the random vector $Y_{s:t} = (Y_1, \ldots, Y_t)$ for arbitrary $s, t \in \mathbb{N}$, $s \leq t$.

Lemma 1. Let $\mathbb{P}_{Y_{s:t}} \in \mathcal{P}((\mathbb{R}^{d_y})^{t-s+1})$ be the probability distribution of $Y_{s:t}$ and λ be the Lebesgue measure on $((\mathbb{R}^{d_y})^{t-s+1}, \mathcal{B}((\mathbb{R}^{d_y})^{t-s+1}))$. Then, for all $0 < s \leq t < \infty$, $\mathbb{P}_{Y_{s:t}}$ is absolutely continuous with respect to λ and its Radon-Nikodym derivative is

$$\frac{d\mathbb{P}_{Y_{s:t}}}{d\lambda}(y_{s:t}) = \Upsilon(y_{s:t}) \triangleq \int_{(\mathbb{R}^{d_x})^{t-s+1}} \prod_{i=s}^{t} g_i(y_i - h(x_i)) \mathbb{P}_{X_{s:t}}(dx_{s:t}), \tag{5}$$

where $\mathbb{P}_{X_{s:t}} \in \mathcal{P}((\mathbb{R}^{d})^{t-s+1})$ is the probability distribution of the random vector $X_{s:t} = (X_s, \ldots, X_t)$.

Remark 2. Note that

$$\mathbb{P}(Y_t \in dy_t \mid X_t = x_t) = g_t(y_t - h(t, x_t))\, dy_t,$$

which explains why the function $g_t^{y_t} : \mathbb{R}^{d_x} \to \mathbb{R}$ defined by

$$g_t^{y_t}(x) = g_t(y_t - h(t, x)), \; x \in \mathbb{R}^{d_x} \tag{6}$$

is commonly referred to as the *likelihood* function.

Since g_i for $i = s, \ldots, t$ are strictly positive, the density of the random vector $(Y_s, \ldots, Y_t)$ is also strictly positive. This condition can be relaxed (i.e. g_i required to be non-negative), however the relaxation requires a more involved theoretical treatment of the particle filter.
The recurrence formula for π_t involves two operations defined on $\mathcal{P}(\mathbb{R}^{d_x})$ – a transformation via the transition kernel K_t and a *projective* product associated with the likelihood function $g_t^{y_t}$ defined as follows:

Definition 3 . Let $p \in \mathcal{P}(\mathbb{R}^{d_x})$ be a probability measure, and let $\varphi \in B(\mathbb{R}^{d_x})$ be a non-negative function such that $p(\varphi) > 0$. The *projective product* $\varphi * p$ is the (set) function $\varphi * p \colon \mathcal{B}(\mathbb{R}^d) \to \mathbb{R}$ defined by

$$\varphi * p(A) \triangleq \frac{\int_A \varphi(x) p(dx)}{p(\varphi)}$$

for any Borel set A.

Exercise 2. Prove that $\varphi * p$ is a probability measure on $\mathcal{B}(\mathbb{R}^d)$.

The projective product $\varphi * p$ is a probability measure which is absolutely continuous with respect to p, whose Radon-Nikodym derivative with respect to p is proportional to φ viz:

$$\frac{d(\varphi * p)}{dp} = c\varphi,$$

where c is the normalizing constant, $c = 1/p(\varphi)$.
The following result gives the recurrence formula for the conditional probability of the signal. The prior and the posterior distributions coincide at time 0,

$$\pi_0 = q_0,$$

since $Y_0 = 0$ i.e. no observations are available at time 0.

Proposition 4. For any fixed path $(y_0, y_1, \ldots, y_t, \ldots)$ the sequence of (non-random) probability measures $(\pi_t)_{t\geq 0}$ satisfy the following recurrence relation

$$\pi_t = g_t^{y_t} * K_t \pi_{t-1}, \ t > 0. \tag{7}$$

The recurrence formula (7) can be re-written in the following expanded way:

$$\pi_{t-1} \mapsto p_t = K_t \pi_{t-1} \mapsto \pi_t = g_t^{Y_t} * p_t, \ t > 0. \tag{8}$$

The first step is called the *prediction* step: it occurs at time t before the arrival of the new observation Y_t. The second step is the *updating* step as it takes into account the new observation Y_t.
The simplicity of the recurrence formula (8) is misleading. A closed formula for the posterior distribution exists only in exceptional cases (the linear/Gaussian filter). The main difficulty resides in the updating step – the projective product is a non-linear transformation involving the computation of the normalizing constant $p_t(g_t^{Y_t})$ which requires an integration over a (possibly) high-dimensional space.

Convergence of Approximations to π_t

As stated above, (8) requires the computation of the predicted probability measure p_t:

$$\pi_{t-1} \longrightarrow p_t \longrightarrow \pi_t.$$

Therefore it is natural to study algorithms which provide recursive approximations for π_t using intermediate approximations for p_t. Denote by $(\pi_t^n)_{n=1}^{\infty}$ and $(p_t^n)_{n=1}^{\infty}$ the approximating sequence for π_t and respectively p_t, which will be assumed to satisfy the following three conditions:

- π_t^n and p_t^n are *random* measures, not necessarily probability measures.
- $p_t^n \neq 0$, $\pi_t^n \neq 0$ (i.e. no approximation should be trivial).
- $p_t^n g_t^{y_t} > 0$ for all $n > 0$, $0 \leq t \leq T$.

Let $\bar{\pi}_t^n$ be defined as a (random) probability measure absolutely continuous with respect to p_t^n for $t \in \mathbb{N}$ and $n \geq 1$ such that

$$\bar{\pi}_t^n = g_t^{y_t} * p_t^n, \tag{9}$$

thus

$$\bar{\pi}_t^n f = \frac{p_t^n(fg^{y_t})}{p_t^n g^{y_t}}. \tag{10}$$

The following theorem gives necessary and sufficient conditions for the convergence of p_t^n to p_t and π_t^n to π_t.

Theorem 5. For all measurable functions f and all $t \in [0, T]$ the convergence

a0. $\lim_{n\to\infty} \mathbb{E}\left[|\pi_t^n f - \pi_t f|\right] = 0.$

b0. $\lim_{n\to\infty} \mathbb{E}\left[|p_t^n f - p_t f|\right] = 0.$

hold true if and only if for all measurable functions f and all $t \in [0, T]$ we have

a1. $\lim_{n\to\infty} \mathbb{E}\left[|\pi_0^n f - \pi_0 f|\right] = 0.$

b1. $\lim_{n\to\infty} \mathbb{E}\left[\left|p_t^n f - K_{t-1}\pi_{t-1}^n f\right|\right] = \lim_{n\to\infty} \mathbb{E}\left[|\pi_t^n f - \bar{\pi}_t^n f|\right] = 0.$

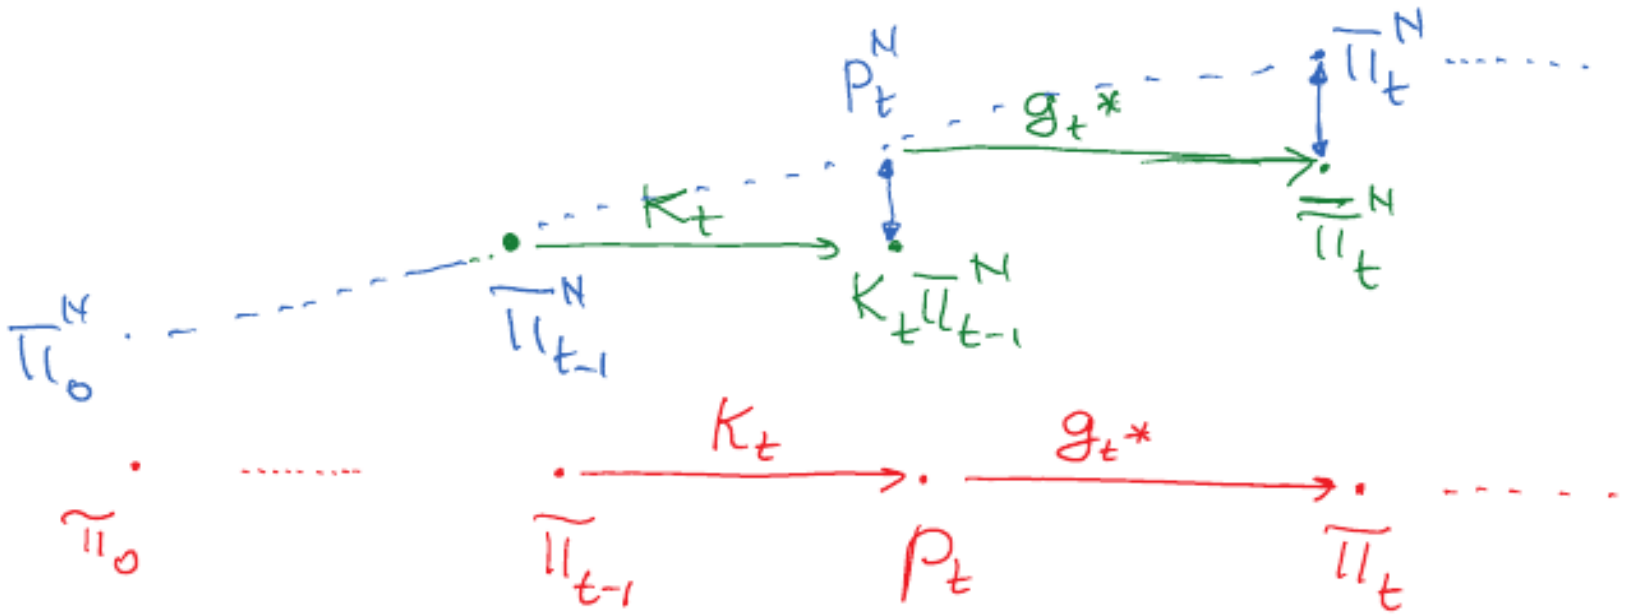

Proof. The necessity of conditions a0. and b0. is proved by induction. The limit a0. follows in the starting case of $t = 0$ from a1. We need to show that if π_{t-1}^n converges in expectation to π_{t-1} and p_t^n converges in expectation to p_t then π_t^n converges in expectation to π_t. Since $p_t = K_t \pi_{t-1}$ by the triangle inequality

$$|p_t^n f - p_t f| \leq |p_t^n f - K_t \pi_{t-1}^n f| + |K_t \pi_{t-1}^n f - K_t \pi_{t-1} f|. \tag{11}$$

The expected value of the first term on the right hand side of (11) converges to zero from b1. Also using **Exercise 1**, $K_{t-1} f \in B(\mathbb{R}^d)$ and $K_{t-1} \pi_{t-1}^n f = \pi_{t-1}^n (K_{t-1} f)$ and $K_{t-1} \pi_{t-1} f = \pi_{t-1}(K_{t-1} f)$ hence

$$\lim_{n \to \infty} \mathbb{E}\left[\left| K_{t-1} \pi_{t-1}^n f - K_{t-1} \pi_{t-1} f \right|\right] = 0.$$

By taking expectation of both sides of (11)

$$\lim_{n \to \infty} \mathbb{E}\left[|p_t^n f - p_t f|\right] = 0, \tag{12}$$

which establishes condition a0.

From (10)

$$\bar{\pi}_t^n f - \pi_t f = \frac{p_t^n(fg_t)}{p_t^n g_t} - \frac{p_t(fg_t)}{p_t g_t}$$
$$= -\frac{p_t^n(fg_t)}{p_t^n g_t}\frac{1}{p_t g_t}(p_t^n g_t - p_t g_t) + \left(\frac{p_t^n(fg_t)}{p_t g_t} - \frac{p_t(fg_t)}{p_t g_t}\right),$$

and as $|p_t^n(fg_t)| \leq \|f\|_\infty p_t^n g_t$,

$$|\bar{\pi}_t^n f - \pi_t f| \leq \frac{\|f\|_\infty}{p_t g_t}|p_t^n g_t - p_t g_t| + \frac{1}{p_t g_t}|p_t^n(fg_t) - p_t(fg_t)|. \tag{13}$$

Therefore

$$\mathbb{E}\left[|\bar{\pi}_t^n f - \pi_t f|\right] \leq \frac{\|f\|_\infty}{p_t g_t}\mathbb{E}\left[|p_t^n g_t - p_t g_t|\right]$$
$$+ \frac{1}{p_t g_t}\mathbb{E}\left[|p_t^n(fg_t) - p_t(fg_t)|\right]. \tag{14}$$

From (12) both terms on the right hand side of (14) converge to zero. Finally,

$$|\pi_t^n f - \pi_t f| \leq |\pi_t^n f - \bar{\pi}_t^n f| + |\bar{\pi}_t^n f - \pi_t f|. \tag{15}$$

As the expected value of the first term on the right hand side of (15) converges to zero using b1. and the expected value of the second term converges to zero using (14), $\lim_{n\to\infty} \mathbb{E}\left[|\pi_t^n f - \pi_t f|\right] = 0$.
For the sufficiency part, assume that conditions a0. and b0. hold. Thus for all $t \geq 0$ and for all $f \in B(\mathbb{R}^d)$,

$$\lim_{n\to\infty} \mathbb{E}\left[|\pi_t^n f - \pi_t f|\right] = \lim_{n\to\infty} \mathbb{E}\left[|p_t^n f - p_t f|\right] = 0.$$

Clearly condition a1. follows as a special case of a0. with $t = 0$. Since $p_t = K_{t-1}\pi_{t-1}$, we have for all $f \in B(\mathbb{R}^d)$,

$$\begin{aligned}\mathbb{E}\left[\left|p_t^n f - K_{t-1}\pi_{t-1}^n f\right|\right] \leq \mathbb{E}\left[|p_t^n f - p_t f|\right] \\ + \mathbb{E}\left[\left|\pi_{t-1}(K_{t-1} f) - \pi_{t-1}^n(K_{t-1} f)\right|\right],\end{aligned} \tag{16}$$

which implies the first limit in b1. From (14),

$$\lim_{n\to\infty} \mathbb{E}\left[|\pi_t f - \bar{\pi}_t^n f|\right] = 0$$

and by the triangle inequality

$$\mathbb{E}\left[|\pi_t^n f - \bar{\pi}_t^n f|\right] \leq \mathbb{E}\left[|\pi_t^n f - \pi_t f|\right] + \mathbb{E}\left[|\pi_t f - \bar{\pi}_t^n f|\right] \tag{17}$$

from which the second limit in b1. follows.

Proposition 6. If there exists a positive constant $p > 1$ such that

$$\mathbb{E}\left[\left|\pi_t^n f - \pi_t f\right|^{2p}\right] \leq \frac{c_f}{n^p}, \tag{18}$$

where c_f is a positive constant depending on the test function f, but independent of n, then, for any $\varepsilon \in (0, 1/2 - 1/(2p))$ there exists a positive random variable $c_{f,\varepsilon}$ almost surely finite such that

$$\left|\pi_t^n f - \pi_t f\right| \leq \frac{c_{f,\varepsilon}}{n^\varepsilon}.$$

In particular $\pi_t^n f$ converges to $\pi_t f$ almost surely.

Proof. Since

$$\mathbb{E}\left[\sum n^{2p\varepsilon}\left|\pi_t^n f - \pi_t f\right|^{2p}\right] \leq \sum n^{2p\varepsilon}\mathbb{E}\left[\left|\pi_t^n f - \pi_t f\right|^{2p}\right] \leq \sum \frac{c_f}{n^{p(1-2\varepsilon)}} < \infty,$$

as $n(1-2\varepsilon) > 1 \Leftrightarrow 1/2 - 1/(2p) > \varepsilon$ it follows that

$$c_{f,\varepsilon} \triangleq \sum n^{2p\varepsilon}\left|\pi_t^n f - \pi_t f\right|^{2p} < \infty.$$

hence the claim.

Moreover if (18) holds for any $f \in \mathcal{M}$ where $\mathcal{M}$ is a countable convergence determining set, then, almost surely, π_t^n converges to π_t in the weak topology. This means that there exists a set $\bar{\Omega} \in \mathcal{F}$ such that $\mathbb{P}(\bar{\Omega}) = 1$ and for any $\omega \in \bar{\Omega}$ the corresponding sequence of probability measures $\pi_t^{n,\omega}$ satisfies

$$\lim_{n\to\infty} \pi_t^{n,\omega}(f) = \pi_t(f),$$

for any $f \in C_b(\mathbb{R}^d)$. This cannot be extended to the convergence for any $f \in B(\mathbb{R}^d)$ (that is to the stronger, so-called convergence in *total variation*).

Exercise 4. Let μ be the uniform measure on the interval $[0, 1]$ and $(\mu_n)_{n\geq 1}$ be the sequence of probability measures $\mu_n = \frac{1}{n}\sum_{i=1}^n \delta_{i/n}$.
i. Show that $(\mu_n)_{n\geq 1}$ converges to μ in the weak topology.
ii. Let $f = 1_{\mathbb{Q}\cap[0,1]}$ be the indicator set of all the rational numbers in $[0, 1]$. Show that $\mu_n(f) \not\to \mu(f)$, hence μ_n does *not* converge to μ in total variation.

Remark 7. Theorem 5 is very natural. It says that we obtain approximations of p_t and π_t for all $t \in [0, T]$ if and only if we start from an approximation of π_0 and then 'follow closely' the recurrence formula (9) for p_t and π_t.

Particle Filters

- In this section we present examples of approximations to the posterior distribution which satisfy the conditions stated in Theorem 5. The algorithms used to produce these approximations are called *particle filters* or *sequential Monte Carlo methods*.
- The algorithms presented below involve the use of a system of n particles which evolve (mutate) according to the law of X.
- After each mutation the system is corrected – each particle is replaced by a random number of particles whose mean is proportional to the likelihood of the position of the particle.
- After imposing some weak restrictions on the offspring distribution of the particles, the empirical measure associated to the particle systems is proven to converge (as n tends to ∞) to the conditional distribution of the signal given the observation.

Denote by π_t^n the approximation to π_t and by p_t^n the approximation to p_t. The particle filter has the following description:

A typical Particle Filter

1. **Initialization** $[t = 0]$.
   For $i = 1, \ldots, n$, sample $x_0^{(i)}$ from π_0.
2. **Iteration** $[t - 1$ to $t]$.
   Let $x_{t-1}^{(i)}$, $i = 1, \ldots, n$ be the positions of the particles at time $t - 1$.
   1. For $i = 1, \ldots, n$, sample $\bar{x}_t^{(i)}$ from $K_{t-1}(x_{t-1}^{(i)}, \cdot)$.
   2. Compute the (normalized) weight $w_t^{(i)} = g_t(\bar{x}_t^{(i)})/(\sum_{j=1}^{n} g_t(\bar{x}_t^{(j)}))$.
   3. Replace each particle by $\xi_t^{(i)}$ offsprings such that $\sum_{i=1}^{n} \xi_t^{(i)} = n$.
   4. Denote the positions of the offspring particles by $x_t^{(i)}$, $i = 1, \ldots, n$.

It follows from the above that the particle filter starts from π_0^n – the empirical measure associated to a set of n random particles of mass $1/n$ whose positions $x_0^{(i)}$ for $i = 1, \ldots, n$ form a sample of size n from π_0

$$\pi_0^n \triangleq \frac{1}{n} \sum_{i=1}^{n} \delta_{x_0^{(i)}}.$$

In general, define π_t^n to be

$$\pi_t^n \triangleq \frac{1}{n} \sum_{i=1}^{n} \delta_{x_t^{(i)}},$$

where $x_t^{(i)}$ for $i = 1, \ldots, n$ are the positions of the particles of mass $1/n$ obtained after the second step of the iteration. Let $\bar{\pi}_t^n$ be the weighted measure

$$\bar{\pi}_t^n \triangleq \sum_{i=1}^{n} w_t^{(i)} \delta_{\bar{x}_t^{(i)}}.$$

We introduce the following σ-algebras

$$\mathcal{F}_t = \sigma(x_s^{(i)}, \bar{x}_s^{(i)}, s \leq t,\ i = 1, \ldots, n)$$
$$\bar{\mathcal{F}}_t = \sigma(x_s^{(i)}, \bar{x}_s^{(i)}, s < t,\ \bar{x}_t^{(i)},\ i = 1, \ldots, n).$$

Obviously $\bar{\mathcal{F}}_t \subset \mathcal{F}_t$ and the (random) probability measures p_t^n and $\bar{\pi}_t^n$ are $\bar{\mathcal{F}}_t$-measurable whilst π_t^n is $\mathcal{F}_t$-measurable for any $t \geq 0$. The random variables $\bar{x}_t^{(i)}$ for $i = 1, \ldots, n$ are chosen to be mutually independent *conditional upon* $\mathcal{F}_{t-1}$.
The iteration uses π_{t-1}^n to obtain π_t^n, but not any of the previous approximations. Following part (a) of the iteration, each particle changes its position according to the transition kernel of the signal. Let p_t^n be the empirical distribution associated with the cloud of particles of mass $1/n$ after part (a) of the iteration

$$p_t^n = \frac{1}{n} \sum_{i=1}^{n} \delta_{\bar{x}_t^{(i)}}.$$

This step of the algorithm is known under the name of *importance sampling* step (popular in the statistics literature) or *mutation* step (inherited from the genetic algorithms literature).

Exercise 5. Prove that $\mathbb{E}\left[p_t^n \mid \mathcal{F}_{t-1}\right] = K_{t-1}^n \pi_{t-1}^n$.

Remark 8. An alternative way to obtain p_t^n from π_{t-1}^n is to sample n times from the measure $K_{t-1}\pi_{t-1}^n$ and define p_t^n to be the empirical measure associated with this sample. Let A_t^n be the conditional covariance matrix of the random vector $\xi_t \triangleq (\xi_t^{(i)})_{i=1}^n$,

$$A_t^n \triangleq \mathbb{E}\left[(\xi_t - nw_t)^\top (\xi_t - nw_t) \mid \bar{\mathcal{F}}_t\right]$$

with entries $(A_t^n)_{ij} = \mathbb{E}\left[\left(\xi_t^{(i)} - nw_t^{(i)}\right)\left(\xi_t^{(j)} - nw_t^{(j)}\right) \,\middle|\, \bar{\mathcal{F}}_t\right]$, where $w_t \triangleq (w_t^{(i)})_{i=1}^n$ is the vector of weights. We assume that the offspring vector $\xi_t = (\xi_t^{(i)})_{i=1}^n$ satisfies the following two conditions:

1. The conditional mean number of offspring is proportional to $w_t^{(i)}$. More precisely
$$\mathbb{E}\left[\xi_t^{(i)} \mid \bar{\mathcal{F}}_t\right] = nw_t^{(i)}. \tag{19}$$
2. There exists a constant c_t, such that
$$q^\top A_t^n q \leq nc_t \tag{20}$$
for any vector $q = \left(q^{(i)}\right)_{i=1}^n \in \mathbb{R}^n$, such that $|q^{(i)}| \leq 1$ for $i = 1, \ldots, n$.

Exercise 6. Prove that the following identity holds

$$\pi_t^n = \frac{1}{n}\sum_{i=1}^{n} \xi_t^{(i)} \delta_{\bar{x}_t^{(i)}},$$

and that $\mathbb{E}[\pi_t^n \mid \bar{\mathcal{F}}_t] = \bar{\pi}_t^n$.

Step (b) of the iteration is called the *selection* step. The particles obtained after the first step of the recursion are multiplied or discarded according to the magnitude of the likelihood weights. In turn the likelihood weights are proportional to the likelihood of the new observation given the corresponding position of the particle. Hence if $nw_t^{(i)}$ is small, fewer offspring are expected than if $nw_t^{(i)}$ is large. Since

$$w_t^{(i)} = \frac{g_t\left(\bar{x}_t^{(i)}\right)}{\frac{1}{n}\sum_{j=1}^{n} g_t\left(\bar{x}_t^{(j)}\right)},$$

$nw_t^{(i)}$ is small when the corresponding value of the likelihood function $g_t(\bar{x}_t^{(i)})$ is smaller than the likelihood function averaged over the positions of all the particles.

In conclusion, the effect of part (b) of the iteration is that it discards particles in unlikely positions and multiplies those in more likely ones. Following Exercise 6, this is done in an *unbiased* manner – the conditional expectation of the approximation after applying the step is equal to the weighted sample obtained after the first step of the recursion. That is, the average of the mass $\xi_t^{(i)}/n$ associated with particle i is equal to $w_t^{(i)}$, the weight of the particle before applying the step.

Exercise 7 Prove that, for all f measurable functions, we have

$$\mathbb{E}\left[\left(\pi_t^n f - \bar{\pi}_t^n f\right)^2\right] \leq \frac{c_t \|f\|_\infty^2}{n}.$$

Exercise 7 implies that the randomness introduced in part (b) of the iteration, as measured by the second moment of $\pi_t^n f - \bar{\pi}_t^n f$ tends to zero with rate given by $1/n$, where n is the number of particles in the system.

Remark 9 Condition (20) is equivalent to

$$q^\top A_t^n q \leq n\bar{c}_t \tag{21}$$

for any vector $q = \left(q^{(i)}\right)_{i=1}^n \in [0,1]^n$, where $\bar{c}_t$ is a fixed constant.

Offspring Distributions

In order to have a complete description of the particle filter we need to specify the offspring distribution. The most popular offspring distribution is the multinomial distribution

$$\xi_t = \text{Multinomial}\left(n, w_t^{(1)}, \ldots, w_t^{(n)}\right)$$

that is

$$\mathbb{P}\left(\xi_t^{(i)} = n^{(i)}, i = 1, \ldots, n\right) = \frac{n!}{\prod_{i=1}^n n^{(i)}!} \prod_{i=1}^n \left(w_t^{(i)}\right)^{n^{(i)}}.$$

The multinomial distribution is the empirical distribution of an n-sample from the distribution $\bar{\pi}_t^n$.

Multinomial Sampling

for $j := 1$ *to* n
Pick $x_t^{(j)}$ *by sampling with replacement from the set of particle positions* $(\overline{x}_t^{(1)}, \overline{x}_t^{(2)}, \ldots, \overline{x}_t^{(n)})$ *according to the probability vector of normalized weights* $(w_t^{(1)}, w_t^{(2)}, \ldots, w_t^{(n)})$.

In other words, if we sample (with replacement) n-times from the population of particles with positions $\bar{x}_t^{(i)}$, $i = 1, \ldots, n$ according to the probability distribution given by the corresponding weights $w_t^{(i)}$, $i = 1, \ldots, n$ and denote by $\xi_t^{(i)}$ the number of times that the particle with position $\bar{x}_t^{(i)}$ is chosen, then $\xi_t = (\xi_t^{(i)})_{i=1}^n$ has the above multinomial distribution.

Lemma 10 If ξ_t has a multinomial distribution then it satisfies the unbiasedness condition, that is

$$\mathbb{E}\left[\xi_t^{(i)} \mid \bar{\mathcal{F}}_t\right] = n w_t^{(i)},$$

for any $i = 1, \ldots, n$. Also ξ_t satisfies condition (20).

Proof. The unbiasedness condition follows immediately from the properties of the multinomial distribution. Also

$$\mathbb{E}\left[\left(\xi_t^{(i)} - n w_t^{(i)}\right)^2 \mid \bar{\mathcal{F}}_t\right] = n w_t^{(i)} \left(1 - w_t^{(i)}\right)$$

$$\mathbb{E}\left[\left(\xi_t^{(i)} - n w_t^{(i)}\right)\left(\xi_t^{(j)} - n w_t^{(j)}\right) \mid \bar{\mathcal{F}}_t\right] = -n w_t^{(i)} w_t^{(j)}, \; i \neq j.$$

Then for all $q = \left(q^{(i)}\right)_{i=1}^{n} \in [-1,1]^n$,

$$
\begin{aligned}
q^\top A_t^n q &= \sum_{i=1}^{n} n w_t^{(i)} \left(1 - w_t^{(i)}\right) \left(q^{(i)}\right)^2 - 2 \sum_{1 \le i < j \le n} n w_t^{(i)} w_t^{(j)} q^{(i)} q^{(j)} \\
&= n \sum_{i=1}^{n} w_t^{(i)} \left(q^{(i)}\right)^2 - n \left(\sum_{i=1}^{n} w_t^{(i)} q^{(i)} \right)^2 \\
&\le n \sum_{i=1}^{n} w_t^{(i)},
\end{aligned}
$$

and since $\sum_{i=1}^{n} w_t^{(i)} = 1$, (20) holds with $c_t = 1$. □

The particle filter with this choice of offspring distribution is called the *Bootstrap Filter* or the *Sampling Importance Resampling* algorithm (SIR algorithm). It was introduced by Gordon Salmond and Smith. Within the context of the Bootstrap Filter, the second step is called the *resampling* step.

Properties of the Bootstrap Filter

- quick and easy to implement
- amenable to parallelization
- suboptimal – the resampling step replaces the (normalized) weights $w_t^{(i)}$ by the random masses $\xi_t^{(i)}/n$, where $\xi_t^{(i)}$ is the number of offsprings of the ith particle. Since ξ_t has a multinomial distribution, $\xi_t^{(i)}$ can take any value between 0 and n.
- even when $w_t^{(i)}$ is high (the position of the ith particle is very likely), the ith particle may have very few offspring or even none at all (albeit with small probability).

If ξ_t is obtained by *residual* sampling, rather than by independent sampling with replacement, then the above disadvantage can be avoided. In this case

$$\xi_t = [nw_t] + \bar{\xi}_t. \tag{22}$$

In (22), $[nw_t]$ is the (row) vector of integer parts of the quantities $nw_t^{(i)}$. That is

$$[nw_t] = \left(\left[nw_t^{(1)} \right], \ldots, \left[nw_t^{(n)} \right] \right),$$

and $\bar{\xi}_t$ has multinomial distribution

$$\bar{\xi}_t = \text{Multinomial}\left(\bar{n}, \bar{w}_t^{(1)}, \ldots, \bar{w}_t^{(n)} \right)$$

where the integer $\bar{n}$ is given by

$$\bar{n} \triangleq n - \sum_{i=1}^{n} \left[nw_t^{(i)} \right] = \sum_{i=1}^{n} \left\{ nw_t^{(i)} \right\}$$

and the weights $\bar{w}_t^{(i)}$ are given by

$$\bar{w}_t^{(i)} \triangleq \frac{\left\{ nw_t^{(i)} \right\}}{\sum_{i=1}^{n} \left\{ nw_t^{(i)} \right\}}.$$

By using residual sampling to obtain ξ_t, we ensure that the original weights $w_t^{(i)}$ are replaced by a random weight which is at least $[nw_t^{(i)}]/n$. This is the closest integer multiple of $1/n$ lower than the actual weight $w_t^{(i)}$. In this way, eliminating particles with likely positions is no longer possible. As long as the corresponding weight is larger than $1/n$, the particle will have *at least* one offspring.

Lemma 11. If ξ_t has distribution given by (22), it satisfies both the unbiasedness condition (19) and condition (20).

Proof. The unbiasedness condition follows from the properties of the multinomial distribution:

$$\begin{aligned}\mathbb{E}\left[\xi_t^{(i)} \mid \bar{\mathcal{F}}_t\right] &= \left[nw_t^{(i)}\right] + \mathbb{E}\left[\bar{\xi}_t^{(i)} \mid \bar{\mathcal{F}}_t\right] \\ &= \left[nw_t^{(i)}\right] + \bar{n}\bar{w}_t^{(i)} \\ &= \left[nw_t^{(i)}\right] + \left\{nw_t^{(i)}\right\} = nw_t^{(i)}.\end{aligned}$$

Also

$$\mathbb{E}\left[\left(\xi_t^{(i)} - nw_t^{(i)}\right)^2 \mid \bar{\mathcal{F}}_t\right] = \mathbb{E}\left[\left(\bar{\xi}_t^{(i)} - \{nw_t^{(i)}\}\right)^2 \mid \bar{\mathcal{F}}_t\right]$$
$$= \bar{n}\bar{w}_t^{(i)}\left(1 - \bar{w}_t^{(i)}\right)$$

and

$$\mathbb{E}\left[\left(\xi_t^{(i)} - nw_t^{(i)}\right)\left(\xi_t^{(j)} - nw_t^{(j)}\right) \mid \bar{\mathcal{F}}_t\right] = -\bar{n}\bar{w}_t^{(i)}\bar{w}_t^{(j)}.$$

Then for all $q = (q^{(i)})_{i=1}^n \in [-1,1]^n$, we have

$$q^\top A_t^n q = \sum_{i=1}^n \bar{n}\bar{w}_t^{(i)}\left(1 - \bar{w}_t^{(i)}\right)\left(q^{(i)}\right)^2 - 2\sum_{1\le i<j\le n} \bar{n}\bar{w}_t^{(i)}\bar{w}_t^{(j)}q^{(i)}q^{(j)}$$
$$= \sum_{i=1}^n \bar{n}\bar{w}_t^{(i)}\left(q^{(i)}\right)^2 - \bar{n}\left(\sum_{i=1}^n \bar{w}_t^{(i)}q^{(i)}\right)^2 \le \sum_{i=1}^n \bar{n}\bar{w}_t^{(i)},$$

and since $\sum_{i=1}^n \bar{n}\bar{w}_t^{(i)} = \sum_{i=1}^n \{nw_t^{(i)}\} < n$, (20) holds with $c_t = 1$. □

The residual sampling is still suboptimal – the correction step now replaces the weight $w_t^{(i)}$ by the deterministic mass $[nw_t^{(i)}]/n$ to which it adds a random mass given by $\bar{\xi}_t^{(i)}/n$, where $\bar{\xi}_t^{(i)}$ can take any value between 0 and $\bar{n}$. This creates a problem for particles with small weights. Even when $w_t^{(i)}$ is small (the position of the ith-particle is very unlikely) it may have a large number of offspring – up to $\bar{n}$ offspring are possible (albeit with small probability). The multinomial distribution also sufffers from this problem.

If ξ_t is obtained by using the branching algorithm described below, then both the above difficulties are eliminated. In this case, the number of offspring $\xi_t^{(i)}$ for each individual particle, has the distribution

$$\xi_t^{(i)} = \begin{cases} \left[nw_t^{(i)}\right] & \text{with probability } 1 - \left\{nw_t^{(i)}\right\} \\ \left[nw_t^{(i)}\right] + 1 & \text{with probability } \left\{nw_t^{(i)}\right\}, \end{cases} \tag{23}$$

whilst $\sum_{i=1}^{n} \xi_t^{(i)}$ remains equal to n.

If the particle has a weight $w_t^{(i)} > 1/n$, then the particle will have offspring. Thus if the corresponding likelihood function $g_t(\bar{x}_t^{(i)})$ is larger than the likelihood averaged over all the existing particles

$$\frac{1}{n}\sum_{j=1}^{n} g_t(\bar{x}_t^{(j)}),$$

then the ith site is selected and the higher the weight $w_t^{(i)}$ the more offspring the ith particle will have. If $w_t^{(i)}$ is less than or equal to $1/n$, the particle will have *at most* one offspring. It will have no offspring with probability $1 - nw_t^{(i)}$, as in this case

$$nw_t^{(i)} = \{nw_t^{(i)}\}.$$

Hence, if $w_t^{(i)} \ll 1/n$, no mass is likely to be assigned to site i. That is, the ith particle is very unlikely and it is eliminated from the sample.

Branching Algorithm

Let u_j, $j = 1, \ldots, n-1$ be $n-1$ mutually independent random variables, uniformly distributed on $[0,1]$, which are independent of all other random variables in the system. The following algorithm is then applied:

$g := n \qquad h := n$
for $i := 1$ *to* $n-1$
  if $\left\{ nw_n^{(i)} \right\} + \left\{ g - nw_n^{(i)} \right\} < 1$ *then*
    if $u_i < 1 - \left(\left\{ nw_n^{(i)} \right\} / \{g\} \right)$ *then*
      $\xi_n^{(i)} := \left[nw_n^{(i)} \right]$
    else
      $\xi_n^{(i)} := \left[nw_n^{(i)} \right] + (h - [g])$
    end if
  else
    if $u_i < 1 - \left(1 - \left\{ nw_n^{(i)} \right\} \right) / (1 - \{g\})$ *then*
      $\xi_n^{(i)} := \left[nw_n^{(i)} \right] + 1$
    else
      $\xi_n^{(i)} := \left[nw_n^{(i)} \right] + (h - [g])$
    end if
  end if
  $g := g - nw_n^{(i)}$
  $h := h - \xi_n^{(i)}$
end for
$\xi_n^{(n)} := h$

We have now n particles with positions

$$(\underbrace{\overline{x}_t^{(1)}, \overline{x}_t^{(1)}, \ldots, \overline{x}_t^{(1)}}_{\xi_n^{(1)}}, \underbrace{\overline{x}_t^{(2)}, \overline{x}_t^{(2)}, \ldots, \overline{x}_t^{(2)}}_{\xi_n^{(2)}}, \ldots) \tag{24}$$

Re-index the positions of the particles as

$$x_t^{(1)}, x_t^{(2)}, \ldots, x_t^{(n)}.$$

The positions of the particles with no offspring will no longer appear among those described by the formula (24).
Some of the properties of the random variables

$$\{\xi_n^{(i)},\ i = 1, \ldots, n\}$$

are given by the following Proposition

Proposition 12 The random variables $\{\xi_n^{(i)},\ i = 1, \ldots, n\}$ have the following properties:

a. $\sum_{i=1}^{n} \xi_n^{(i)} = n$.

b. For any $i = 1, \ldots, n$ we have $\mathbb{E}[\xi_n^{(i)}] = nw_n^{(i)}$.

c. For any $i = 1, \ldots, n$, $\xi_n^{(i)}$ has minimal variance, specifically

$$\mathbb{E}[(\xi_n^{(i)} - nw_n^{(i)})^2] = \{nw_n^{(i)}\}(1 - \{nw_n^{(i)}\}).$$

d. For $1 \leq i < j \leq n$, the random variables $\xi_n^{(i)}$ and $\xi_n^{(j)}$ are negatively correlated. That is

$$\mathbb{E}[(\xi_n^{(i)} - nw_n^{(i)})(\xi_n^{(j)} - nw_n^{(j)})] \leq 0.$$

Properties

- If ξ_t is obtained as above, then it is optimal in the sense that, for any $i = 1, \ldots, n$, $\xi_t^{(i)}$ has the *smallest possible variance* amongst all integer-valued random variables with the given mean $nw_t^{(i)}$.
- The algorithm ensures that minimal randomness, as measured by the variance of the mass allocated to individual sites, is introduce to the system.
- The minimal variance property for the distribution produced by any tree based branching algorithm holds true not only for individual sites but also for all groups of sites corresponding to a node of the building binary tree.
- A second optimality property of this distribution is that it has the minimal relative entropy with respect to the measure $\bar{\pi}_t$ which it replaces in the class of all empirical distribution of n particles of mass $1/n$.

Lemma 13 If ξ_t is produced by the algorithm described above, it satisfies both unbiasedness condition (19) and condition (20).

Proof. The unbiasedness condition immediately follows from Proposition 12

$$\mathbb{E}\left[\xi_t^{(i)} \mid \bar{\mathcal{F}}_t\right] = \left[nw_t^{(i)}\right]\left(1 - \left\{nw_t^{(i)}\right\}\right) + \left(\left[nw_t^{(i)}\right] + 1\right)\left\{nw_t^{(i)}\right\} = nw_t^{(i)}.$$

Also

$$\mathbb{E}\left[\left(\xi_t^{(i)} - nw_t^{(i)}\right)^2 \mid \bar{\mathcal{F}}_t\right] = \left\{nw_t^{(i)}\right\}\left(1 - \left\{nw_t^{(i)}\right\}\right),$$

and from Proposition 12 part (d),

$$\mathbb{E}\left[\left(\xi_t^{(i)} - nw_t^{(i)}\right)\left(\xi_t^{(j)} - nw_t^{(j)}\right) \mid \bar{\mathcal{F}}_t\right] \leq 0.$$

Then for all $q = (q^{(i)})_{i=1}^n = [0,1]^n$, we have

$$q^\top A_t^n q \leq \sum_{i=1}^n \left\{nw_t^{(i)}\right\}\left(1 - \left\{nw_t^{(i)}\right\}\right),$$

and since $\{nw_t^{(i)}\}(1 - \{nw_t^{(i)}\}) < 1/4$, following Lemma **??**, condition (20) holds with $c_t = 1/4$. □

There exists another algorithm that satisfies the same minimal variance property of the branching algorithm described above. It was introduced by Carpenter, Clifford and Fearnhead in the context of particle approximations. The method had appeared earlier in the field of genetic algorithms and it is known under the name of *stochastic universal sampling* (see Baker and Whitley). However the offspring distribution generated by this method *does not* satisfy condition (20) and the convergence of the particle filter with this method is still an open question.

All offspring distributions presented above leave the total number of particles constant and satisfy (20). However, the condition that the total number of particles does not change is not essential.

One can choose the individual offspring numbers $\xi_t^{(i)}$ to be mutually independent given $\bar{\mathcal{F}}_t$. As alternatives for the distribution of the integer-valued random variables $\xi_t^{(i)}$ the following can be used:

1. $\xi_t^{(i)} = B(n, w_t^{(i)})$, that is, $\xi_t^{(i)}$ are binomially distributed with parameters $(n, w_t^{(i)})$.
2. $\xi_t^{(i)} = P(nw_t^{(i)})$, that is, $\xi_t^{(i)}$ are Poisson distributed with parameters $nw_t^{(i)}$.
3. $\xi_t^{(i)}$ are Bernoulli distributed with distribution given by (23).

Exercise 8. Show that if the individual offspring numbers $\xi_t^{(i)}$ are mutually independent given $\bar{\mathcal{F}}_t$ and have any of the three distributions described above, then ξ_t satisfies both unbiasedness condition and condition (20).
The Bernoulli distribution is the optimal choice for independent offspring distributions. Since $\sum_{i=1}^{n} \xi_t^{(i)}$ is no longer equal to n, the approximating measure π_t^n is no longer a probability measure. However, following the unbiasedness condition (19) and condition (20), the total mass $\pi_t^n(\mathbf{1})$ of the approximating measure is a martingale which satisfies, for any $t \in [0, T]$,

$$\mathbb{E}\left[\left(\pi_t^n(\mathbf{1}) - 1\right)^2\right] \leq \frac{c}{n},$$

where $c = c(T)$ is a constant independent of n. This implies that for large n the mass oscillations becomes very small. Indeed, by Chebyshev's inequality

$$\mathbb{P}\left(\left|\pi_t^n(\mathbf{1}) - 1\right| \geq \varepsilon\right) \leq \frac{c}{n\varepsilon^2}.$$

- Hence, having a non-constant number of particles does not necessarily lead to instability.
- The oscillations in the number of particles can in themselves constitute an indicator of the convergence of the algorithm.
- Such an offspring distribution with independent individual offspring numbers is easy to implement and saves computational effort.
- Theorem 5 can be used in order to prove the convergence of any algorithm based on such offspring distributions.

Convergence of SMC methods

Exercise 9. Prove that

$$\lim_{n\to\infty} \mathbb{E}\left[|\pi_0^n f - \pi_0 f|\right] = 0.$$

Theorem 14. Let $(p_t^n)_{n=1}^{\infty}$ and $(\pi_t^n)_{n=1}^{\infty}$ be the measure-valued sequences produced by the class of algorithms described above. Then, for all $0 \leq t \leq T$, we have

$$\lim_{n\to\infty} \mathbb{E}\left[|\pi_t^n f - \pi_t f|\right] = \lim_{n\to\infty} \mathbb{E}\left[|p_t^n f - p_t f|\right] = 0,$$

for all $f \in B(\mathbb{R}^d)$.

Proof. We apply Theorem 5. Since a1. holds as a consequence of Exercise 9, it is only necessary to verify condition b1. From Exercise 5, $\mathbb{E}\left[p_t^n f \mid \mathcal{F}_t\right] = \pi_{t-1}^n(K_{t-1} f)$ and using the independence of the sample $\{\bar{x}_t^{(i)}\}_{i=1}^n$ conditional on $\mathcal{F}_{t-1}$,

$$
\begin{aligned}
\mathbb{E}\left[\left(p_t^n f - \pi_{t-1}^n(K_{t-1}f)\right)^2 \mid \mathcal{F}_{t-1}\right] &= \frac{1}{n^2}\mathbb{E}\left[\left(\sum_{i=1}^n f\left(\bar{x}_t^{(i)}\right) - K_{t-1}f\left(x_{t-1}^{(i)}\right)\right)^2 \middle| \mathcal{F}_{t-}\right. \\
&= \frac{1}{n^2}\sum_{i=1}^n \mathbb{E}\left[\left(f\left(\bar{x}_t^{(i)}\right)\right)^2 \mid \mathcal{F}_{t-1}\right] \\
&\quad -\frac{1}{n^2}\sum_{i=1}^n \left(\mathbb{E}\left[K_{t-1}f\left(x_{t-1}^{(i)}\right) \mid \mathcal{F}_{t-1}\right]\right)^2 \\
&= \frac{1}{n}\pi_{t-1}^n\left(K_{t-1}f^2 - (K_{t-1}f)^2\right).
\end{aligned}
$$

Therefore $\mathbb{E}[(p_t^n f - \pi_{t-1}^n K_{t-1} f)^2] \leq \|f\|_\infty^2/n$ and the first limit in b1. is satisfied. The second limit in b1. follows from Exercise 7. □

Corollary 15. For all $0 \leq t \leq T$, there exists a constant k_t such that

$$\mathbb{E}\left[(\pi_t^n f - \pi_t f)^2\right] \leq \frac{k_t \|f\|_\infty^2}{n}, \tag{25}$$

for all $f \in B(\mathbb{R}^d)$.

Proof. We proceed by induction. Since $\{x_0^{(i)}, i = 1, \ldots, n\}$ is an n-independent sample from π_0,

$$\mathbb{E}\left[(\pi_0^n f - \pi_0 f)^2\right] \leq \frac{\|f\|_\infty^2}{n},$$

hence by Jensen's inequality (25) is true for $t = 0$ with $k_0 = 1$. Now assume that (25) holds at time $t - 1$. Then

$$\mathbb{E}\left[(\pi_{t-1}^n(K_{t-1} f) - \pi_{t-1}(K_{t-1} f))^2\right] \leq \frac{k_{t-1}\|K_{t-1} f\|_\infty^2}{n} \leq \frac{k_{t-1}\|f\|_\infty^2}{n}. \tag{26}$$

Also from the proof of Theorem 14

$$\mathbb{E}\left[\left(p_t^n f - \pi_{t-1}^n K_{t-1} f\right)^2\right] \leq \frac{\|f\|_\infty^2}{n}. \tag{27}$$

By using inequality (11) and the triangle inequality for the L_2-norm

$$\mathbb{E}\left[(p_t^n f - p_t f)^2\right] \leq \frac{\hat{k}_t \|f\|_\infty^2}{n}. \tag{28}$$

where $\hat{k}_t = (\sqrt{k_{t-1}} + 1)^2$. In turn, (28) and (13) imply that

$$\mathbb{E}\left[(\bar{\pi}_t^n f - \pi_t f)^2\right] \leq \frac{\bar{k}_t \|f\|_\infty^2}{n}. \tag{29}$$

where $\bar{k}_t = 4\hat{k}_t \|g_t\|_\infty^2/(p_t g_t)^2$. From Exercise 7.

$$\mathbb{E}\left[(\pi_t^n f - \bar{\pi}_t^n f)^2\right] \leq \frac{c_t \|f\|_\infty^2}{n}, \tag{30}$$

where c_t is the constant appearing in (20). Finally from (29), (30) and the triangle inequality (15), (25) holds with $k_t = (\sqrt{c_t} + \sqrt{\bar{k}_t})^2$. This completes the induction step. □

Condition (20) is essential in establishing the above rate of convergence. A more general condition than (20) is possible, for example, that there exists $\alpha > 0$ such that

$$q^\top A_t^n q \leq n^\alpha c_t \tag{31}$$

for any $q \in [-1, 1]^n$. In this case, inequality (30) would become

$$\mathbb{E}[(\pi_t^n f - \bar{\pi}_t^n f)^2] \leq \frac{c_t \|f\|_\infty^2}{n^{2-\alpha}}.$$

Hence the overall rate of convergence would take the form $\mathbb{E}\left[(\pi_t^n f - \pi_t f)^2\right] \leq \frac{k_t \|f\|_\infty^2}{n^{\max(2-\alpha,1)}}$ for all $f \in B(\mathbb{R}^d)$. Hence if $\alpha > 1$ we will see a deterioration in the overall rate of convergence. On the other hand, if $\alpha < 1$ no improvement in the rate of convergence is obtained as the error in all the other steps of the particle filter remains of order $1/n$. So $\alpha = 1$ is the most suitable choice for condition (20).

Concluding remarks

- Theorem 5 provides an efficient technique for proving convergence of particle algorithms. The necessary and sufficient conditions stated in the theorem are natural and easy to verify.
- They can be applied when the algorithms studied provide both π_t^n (the approximation to π_t) and also p_t^n (the intermediate approximation to p_t). Algorithms are possible where π_t^n is obtained from π_{t-1}^n without using the approximation for p_t. In other words one can perform the mutation step using a different transition from that of the signal. In the statistics literature, the transition kernel K_t is usually called the *importance distribution*. Should a kernel (or importance distribution) $\bar{K}_t$ be used which is different from that of the signal K_t, the form of the weights appearing in the selection step of the particle filter must be changed. The results presented above then apply for p_t now given by $\bar{K}_{t-1}\pi_{t-1}$ and the weighted measure $\bar{\pi}_t^n$ defined in (9) given by $\bar{\pi}_t^n = \sum_{i=1}^{n} \bar{w}_t^{(i)} \delta_{\bar{x}_t^{(i)}}$, where $\bar{w}_t^{(i)}$ are the new weights.

- The randomness introduced in the system at each selection step must be kept to a minimum as it affects the rate of convergence of the algorithm. Therefore one should not apply the selection step after every new observation arrives. Assume that the information received from the observation is 'bad' (i.e. the signal to noise ratio is small). Because of this, the likelihood function is close to being constant and the corresponding weights are all (roughly) equal, $\bar{w}_t^{(i)} \simeq 1/n$. In other words, the observation is uninformative – it cannot distinguish between different sites and all particles are equally likely. In this case no selection procedure needs to be performed. The observation is stored in the weights of the approximation $\bar{\pi}_t^n$ and carried forward to the next step. If a correction procedure is nevertheless performed and ξ_t has a minimal variance distribution, all particles will have a single offspring 'most of the time'. In other words the system remains largely unchanged with high probability. However with small probability, the ith particle might have no offspring (if $\bar{w}_t^{(i)} < 1/n$) or two offspring (if $\bar{w}_t^{(i)} > 1/n$). Hence randomness still enters the system and this can affect the convergence rates. If ξ_t does not have a minimal variance distribution, the amount of randomness is even higher. It remains an open question as to when and how often one should use the selection procedure.

High dimensional problems are harder than their low dimensional counterparts. Example:

Consider

- $\Pi_1 = \mathcal{N}((0,\ldots,0), I_d)$ (mean $(0,\ldots,0)$ and covariance matrix I_d).
- $\Pi_2 = \mathcal{N}((1,\ldots,1), I_d)$ (mean $(1,\ldots,1)$ and covariance matrix I_d).
- $d(\Pi_1,\Pi_2)_{TV} = 2\,\mathbb{P}[\,|X| \leq d/2\,]$, $X \sim N(0,1)$.
- as d increases, the two measures get further and further apart, becoming singular w.r.t. each other exponentially fast.
- it becomes increasingly harder to construct a sample from Π_2 by using a proposal from Π_1.

Solution: The problem of 'moving' from Π_1 to Π_2 is equivalent to that of moving from a standard normal distribution $\mathcal{N}(0,1)$ to a normal distribution $\mathcal{N}(d,1)$ (the total variation distance between $\mathcal{N}(0,1)$ and $\mathcal{N}(d,1)$ is the same as that between Π_1 and Π_2). Rather than jumping from $\mathcal{N}(0,1)$ to $\mathcal{N}(d,1)$ in one step we get there in d steps: at each step moving from $\mathcal{N}(k-1,1)$ to $\mathcal{N}(k,1)$ for index $k = 1,2,\ldots,d$. This algorithm can be immediately transferred to the corresponding multidimensional set-up.

2D Stochastic Navier-Stokes equation on the torus $\mathbb{T}^2 \triangleq [0, L) \times [0, L)$ with periodic boundary conditions:

$$\frac{\partial u}{\partial t} - \nu \Delta u + u \cdot \nabla u + \nabla p = f + W(t, x) \qquad \text{for all } (x, t) \in \mathbb{T}^2 \times (0, \infty), \quad (32)$$

$$\nabla \cdot u = 0 \qquad \text{for all } (x, t) \in \mathbb{T}^2 \times (0, \infty),$$

$$u(x, 0) = u_0(x) \qquad \text{for all } x \in \mathbb{T}^2.$$

- $u : \mathbb{T}^2 \times [0, \infty) \to \mathbb{R}^2$ - the velocity
- $p : \mathbb{T}^2 \times [0, \infty) \to \mathbb{R}^2$ - the pressure
- $f : \mathbb{T}^2 \to \mathbb{R}^2$ - the forcing
- $W(t, x)$ - noise

$$u = \sum_{k \in \mathbb{Z}^2 \setminus \{0\}} u_k(t)\psi_k(x).$$

where

$$\psi_k(x) \triangleq \frac{k^\perp}{|k|} \exp\left(\frac{2\pi i k \cdot x}{L}\right) \quad k = (k_1, k_2)^\top \in \mathbb{Z}^2 \setminus \{0\} \quad k^\perp = (k_2, -k_1)^\top.$$

The equations for the modes:

$$du_k(t) = \left(-\nu\lambda_k u_k(t) - \alpha_k^{l,j} \sum_{l+j=k} u_l(t)u_j(t) + f_k\right) dt + \varepsilon_k dW_t^k.$$

We approximate the modes $u_k(t)$, with $\tilde{u}_k(t)$ for each $k \in \mathbb{Z} \setminus \{0\}$ with $|2\pi k|^2 < \lambda L^2$:

$$d\tilde{u}_k(t) = \left(-\nu\lambda_k \tilde{u}_k(t) - \alpha_k^{l,j} \sum_{\Gamma} \tilde{u}_l(t)\tilde{u}_j(t) + f_k\right) dt + \varepsilon_k dW_t^k; \tag{33}$$

where the set $\Gamma \triangleq \left\{(l,j) \middle| l + j = k \text{ and } |2\pi l|^2 < \lambda L^2 \text{ and } |2\pi j|^2 < \lambda L^2\right\}$.

Model parameters

- we use $k1, k2 = -32, ..., 0, ...32$ (i.e. a 64^2 grid for the discrete fourier components).
- Smoothing problem approximate $p(x_0|y_{1:5})$ where each y_i is a $4x4$ grid on the torus and

$$y_i(j) = u(x_j, t_i) + N(0, 0.2).$$

- the dynamics are initialised by a random sample from the prior $N(0, \delta A^{\alpha})$
- for the prior, $\delta = 5$ and $\alpha = 2.2$.
- torus size is 2π.
- forcing is $\nabla cos(\kappa \cdot x)$ with $\kappa = (1, 1)$ for the stationary regime and $\kappa = (5, 5)$ for the chaotic regime
- ν is $1/50$ for chaotic and $1/10$ for stationary

MCMC plot: computational cost involving 9×10^5 calls of the PDE solver, with the slow mixing need more than 9 days for a decent but not super-reliable answer.

SMC plots: *parallel computing* computation cost involving 14×10^5 calls of the PDE solver $N = 1000$ particles need 7.4 hours.

Numerics done by N. Kantas (Imperial College London).

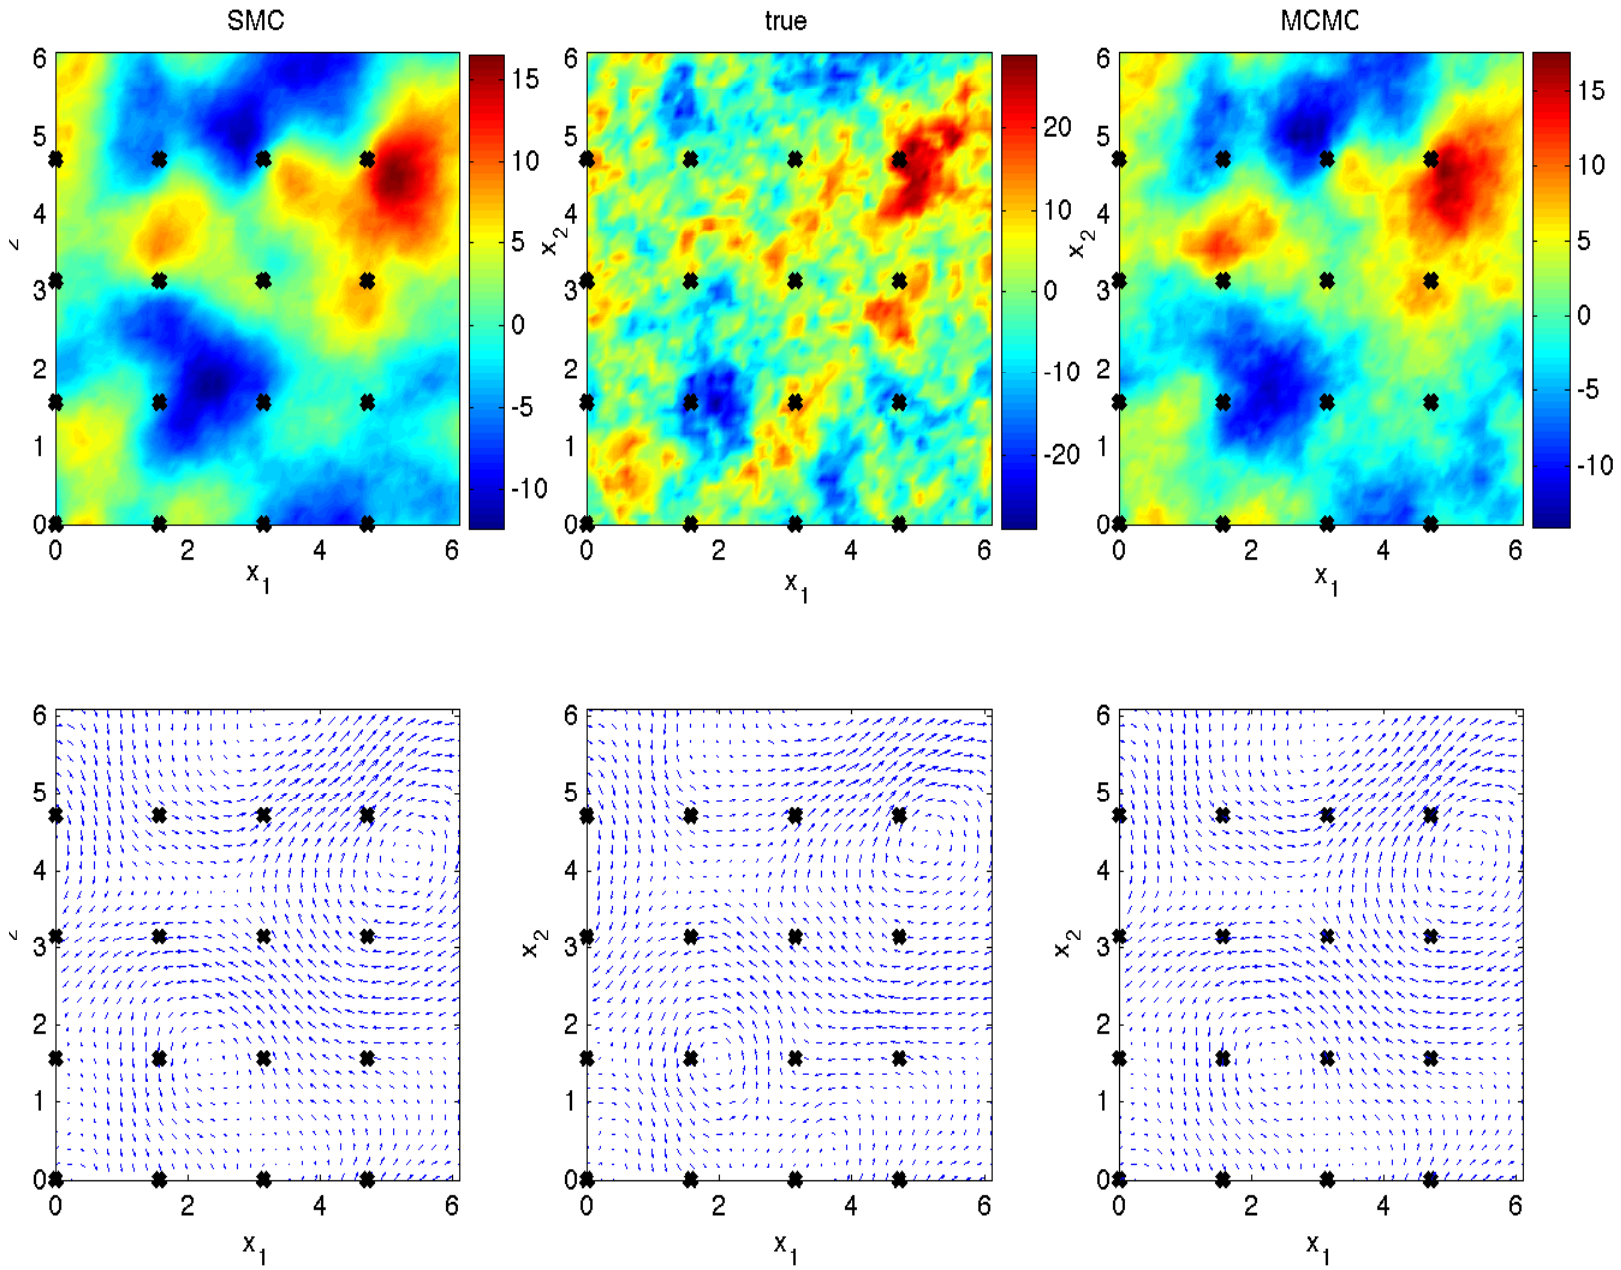
SMC
true
MCMC
x_1
x_2

Topics for future work:

- systematic exploration of DA/stochastic filtering for multi-scale processes
- mathematical theory for DA/stochastic filtering in an infinite-dimensional state space
- stochastic filtering/DA under systematic model and representation errors (e.g. numerical approximation or parametrization errors)
- combined state and parameter estimation
- observation networks, quality of data, etc.
- proposal steps and nonlinear ensemble transform filters