

A nested filtering scheme for Bayesian model inference and tracking

Sara Pérez Vieites

Department of Signal Theory & Communications,
Universidad Carlos III de Madrid.
spvieites@tsc.uc3m.es

Supervisor: Joaquín Míguez

December 5th, 2017

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work



State-of-the-art Methods

- Very-large-scale stochastic dynamic models of real world phenomena, where we estimate
 - dynamic variables and
 - unknown static parameters.

⇒ Conventional prediction and estimation methods are not well suited.
- State-of-the-art algorithms:
 1. Bayesian methods (SMC² [Chopin et al, 2012], pMCMC [Andrieu et al, 2010])
 - approximate the posterior distribution of the unknown variables and parameters, but
 - they are non-recursive.
 2. Maximum likelihood estimators
 - provide point estimates only
 - recursive, but no measure of uncertainty



Contributions

- A purely recursive generic nested filtering scheme is introduced.
 1. The nested particle filter [Crisan and Miguez, 2016] is a particular case, which infers the posterior distribution of variables and parameters
 - using a bank of particle filters.
 - it is computationally costly.
 2. Therefore, a new class of nested hybrid filters is proposed, using
 - a bank of Gaussian filters such as EKF or EnKF, and
 - reducing running times without significant changes in accuracy.
- A stochastic two-scale Lorenz 96 model is used in computer simulations.
 - ⇒ It is a benchmark system for meteorology.

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

State-space model

- The dynamical behaviour of the system state can be modeled by sets of nonlinear stochastic differential equations (SDEs), like:

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, \boldsymbol{\theta})dt + dW \quad (1)$$

where $\mathbf{f}(\tilde{\mathbf{x}})$ is a nonlinear function parametrized by $\boldsymbol{\theta}$ and dW is a brownian process.

State-space model

- The value of the system state can be approximated applying any sort of discretization method as

$$\tilde{\mathbf{x}}_k = \tilde{\mathbf{x}}_{k-1} + h\bar{\mathbf{f}}_\sigma(\tilde{\mathbf{x}}_{k-1}, \boldsymbol{\theta}, \mathbf{v}_k) \quad (2)$$

where

- $\bar{\mathbf{f}}_\sigma(\tilde{\mathbf{x}}_{k-1}, \boldsymbol{\theta}, \mathbf{v}_k)$ is an estimate of the vector of time derivatives
- $h > 0$ is a time-discretisation step with $k = 0, 1, \dots$ and
- $\sigma \geq 0$ is a parameter that controls the power of the perturbations, that needs to be small enough to preserve the underlying dynamics of the system

Observations

States, $\tilde{\mathbf{x}}_k$, and unknown parameters, $\boldsymbol{\theta}$, are estimated from a sequence of observation vectors, modelled as

$$\tilde{\mathbf{y}}_{kT} = \mathbf{g}_{\sigma_o}(\tilde{\mathbf{x}}_{kT}, \boldsymbol{\theta}, \tilde{\mathbf{u}}_{kT}), \quad k = 1, 2, \dots, \quad T \geq 1, \quad (3)$$

where

- $\mathbf{g}_{\sigma_o} : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$, being the observation vectors of dimension $d_y \leq d_x$,
- T is the discrete observation period and
- $\tilde{\mathbf{u}}_k$ is a sequence of zero-mean independent vectors representing observational noise, where power is scaled by a known factor $\sigma_o > 0$.

Dynamical Model

- As we only consider observations **every T discrete time steps**, it is convenient to rewrite the dynamic model in the same time scale. We work with the pair of random sequences $\mathbf{x}_n := \tilde{\mathbf{x}}_{nT}$ and $\mathbf{y}_n := \tilde{\mathbf{y}}_{nT}$, as

$$\mathbf{y}_n = \mathbf{g}_{\sigma_o}(\mathbf{x}_n, \boldsymbol{\theta}, \mathbf{u}_n), \quad (4)$$

$$\mathbf{x}_n = \bar{\mathbf{F}}_{T,\sigma}(\mathbf{x}_{n-1}, \boldsymbol{\theta}, \mathbf{v}_n), \quad n = 1, 2, \dots \quad (5)$$

- Here $\bar{\mathbf{F}}$ represents the transformation from $\mathbf{x}_{n-1} := \tilde{\mathbf{x}}_{(n-1)T}$ to $\mathbf{x}_n := \tilde{\mathbf{x}}_{nT}$ in T steps as

$$\begin{aligned} \tilde{\mathbf{x}}_{(n-1)T} &= \mathbf{x}_{n-1} \\ \tilde{\mathbf{x}}_{(n-1)T+1} &= \tilde{\mathbf{x}}_{(n-1)T} + \bar{\mathbf{f}}_{\sigma}(\tilde{\mathbf{x}}_{(n-1)T}, \boldsymbol{\theta}, \tilde{\mathbf{v}}_{(n-1)T+1}) \\ &\vdots \\ \tilde{\mathbf{x}}_{nT-1} &= \tilde{\mathbf{x}}_{nT-2} + \bar{\mathbf{f}}_{\sigma}(\tilde{\mathbf{x}}_{nT-2}, \boldsymbol{\theta}, \tilde{\mathbf{v}}_{nT-1}) \\ \mathbf{x}_n &= \tilde{\mathbf{x}}_{nT-1} + \bar{\mathbf{f}}_{\sigma}(\tilde{\mathbf{x}}_{nT-1}, \boldsymbol{\theta}, \tilde{\mathbf{v}}_{nT}) \end{aligned} \quad (6)$$

Objective

- The aim of this work is to estimate the parameters θ and its posterior distribution $p(\theta|\mathbf{y})$.
- We combine Monte Carlo with other auxiliary filters for that purpose.

Nested filtering

- The following (naive) algorithm yields a weighted Monte Carlo approximation of $p(\theta|\mathbf{y}_{1:n})d\theta$ at each time n :

- Draw N i.i.d. samples θ_n^i , $i = 1, 2, \dots, N$, from $p(\theta|\mathbf{y}_{1:n-1})$.
- Compute importance weights,

$$\tilde{w}_n^i = p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1}), \quad i = 1, \dots, N, \quad (7)$$

and normalise them. We obtain the IS estimate

$$p(\theta|\mathbf{y}_{1:n})d\theta = \sum_{i=1}^N w_n^i \delta_{\theta_n^i}(d\theta).$$

- This weighted Monte Carlo approximation is not practical because
 - it is not possible to draw from $p(\theta|\mathbf{y}_{1:n})$, at least exactly, and
 - the likelihood $u_n(\theta_n^i) := p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1})$ cannot be evaluated exactly either.
- We can obtain an estimate of $u_n(\theta)$ if we approximate first the predictive measure $p(x_n|\theta, \mathbf{y}_{1:n-1})dx_n$ since it can be rewritten as

$$u_n(\theta) = \int p(\mathbf{y}_n|x_n, \theta)p(x_n|\theta, \mathbf{y}_{1:n-1})d(x_n) \quad (8)$$

Nested filtering

- The following (naive) algorithm yields a weighted Monte Carlo approximation of $p(\theta|\mathbf{y}_{1:n})d\theta$ at each time n :

1. Draw N i.i.d. samples θ_n^i , $i = 1, 2, \dots, N$, from $p(\theta|\mathbf{y}_{1:n-1})$.
2. Compute importance weights,

$$\tilde{w}_n^i = p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1}), \quad i = 1, \dots, N, \quad (7)$$

and normalise them. We obtain the IS estimate

$$p(\theta|\mathbf{y}_{1:n})d\theta = \sum_{i=1}^N w_n^i \delta_{\theta_n^i}(d\theta).$$

- This weighted Monte Carlo approximation is not practical because
 - it is not possible to draw from $p(\theta|\mathbf{y}_{1:n})$, at least exactly, and
 - the likelihood $u_n(\theta_n^i) := p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1})$ cannot be evaluated exactly either.
- We can obtain an estimate of $u_n(\theta)$ if we approximate first the predictive measure $p(x_n|\theta, \mathbf{y}_{1:n-1})dx_n$ since it can be rewritten as

$$u_n(\theta) = \int p(\mathbf{y}_n|x_n, \theta)p(x_n|\theta, \mathbf{y}_{1:n-1})d(x_n) \quad (8)$$

Nested filtering

- The following (naive) algorithm yields a weighted Monte Carlo approximation of $p(\theta|\mathbf{y}_{1:n})d\theta$ at each time n :

1. Draw N i.i.d. samples θ_n^i , $i = 1, 2, \dots, N$, from $p(\theta|\mathbf{y}_{1:n-1})$.
2. Compute importance weights,

$$\tilde{w}_n^i = p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1}), \quad i = 1, \dots, N, \quad (7)$$

and normalise them. We obtain the IS estimate

$$p(\theta|\mathbf{y}_{1:n})d\theta = \sum_{i=1}^N w_n^i \delta_{\theta_n^i}(d\theta).$$

- This weighted Monte Carlo approximation is not practical because
 - it is not possible to draw from $p(\theta|\mathbf{y}_{1:n})$, at least exactly, and
 - the likelihood $u_n(\theta_n^i) := p(\mathbf{y}_n|\theta_n^i, \mathbf{y}_{1:n-1})$ cannot be evaluated exactly either.
- We can obtain an estimate of $u_n(\theta)$ if we approximate first the predictive measure $p(\mathbf{x}_n|\theta, \mathbf{y}_{1:n-1})d\mathbf{x}_n$ since it can be rewritten as

$$u_n(\theta) = \int p(\mathbf{y}_n|\mathbf{x}_n, \theta)p(\mathbf{x}_n|\theta, \mathbf{y}_{1:n-1})d(\mathbf{x}_n) \quad (8)$$

General Nested Filter

Algorithm

1. Initialisation

Draw $\theta_0^{(i)}, i = 1, \dots, N$, i.i.d. samples from $p(\theta)d\theta$.

2. Recursive step

2.1 For $i = 1, \dots, N$:

2.1.1 Draw $\bar{\theta}_n^{(i)}$ from $\kappa_N(d\theta|\theta_{n-1}^i)$.

2.1.2 Approximate $p(x_n|\bar{\theta}_n^i, y_{1:n-1})dx_n$.

2.1.3 Use $p(x_n|\bar{\theta}_n^i, y_{1:n-1})dx_n$ to

compute $\hat{u}_n(\bar{\theta}_n^i)$ and let

$w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

2.2 Resample to obtain the set $\{\theta_{i=1}^N\}$ and the approximation

$$\hat{p}(d\theta|y_{1:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}(d\theta).$$

General Nested Filter

Algorithm

1. Initialisation

Draw $\theta_0^{(i)}, i = 1, \dots, N$, i.i.d. samples from $p(\theta)d\theta$.

2. Recursive step

2.1 For $i = 1, \dots, N$:

2.1.1 Draw $\bar{\theta}_n^{(i)}$ from $\kappa_N(d\theta|\theta_{n-1}^i)$.

2.1.2 Approximate $p(x_n|\bar{\theta}_n^i, y_{1:n-1})dx_n$.

2.1.3 Use $p(x_n|\bar{\theta}_n^i, y_{1:n-1})dx_n$ to

compute $\hat{u}_n(\bar{\theta}_n^i)$ and let

$w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

2.2 Resample to obtain the set $\{\theta_{i=1}^N\}$ and the approximation

$$\hat{p}(d\theta|y_{1:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}(d\theta).$$

Sampling $\bar{\theta}_n^{(i)}$

General Nested Filter

Algorithm

1. Initialisation

Draw $\theta_0^{(i)}, i = 1, \dots, N$, i.i.d. samples from $p(\theta)d\theta$.

2. Recursive step

2.1 For $i = 1, \dots, N$:

2.1.1 Draw $\bar{\theta}_n^{(i)}$ from $\kappa_N(d\theta|\theta_{n-1}^i)$.

2.1.2 Approximate $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$.

2.1.3 Use $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$ to compute $\hat{u}_n(\bar{\theta}_n^i)$ and let $w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

2.2 Resample to obtain the set $\{\theta_{i=1}^N\}$ and the approximation

$$\hat{p}(d\theta|\mathbf{y}_{1:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}(d\theta).$$

Sampling $\bar{\theta}_n^{(i)}$

↓ Prediction

$$\hat{p}(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$$

General Nested Filter

Algorithm

1. Initialisation

Draw $\theta_0^{(i)}$, $i = 1, \dots, N$, i.i.d. samples from $p(\theta)d\theta$.

2. Recursive step

2.1 For $i = 1, \dots, N$:


2.1.1 Draw $\bar{\theta}_n^{(i)}$ from $\kappa_N(d\theta|\theta_{n-1}^i)$.

2.1.2 Approximate $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$.

2.1.3 Use $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$ to compute $\hat{u}_n(\bar{\theta}_n^i)$ and let $w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

2.2 Resample to obtain the set $\{\theta_{i=1}^N\}$ and the approximation

$$\hat{p}(d\theta|\mathbf{y}_{1:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}(d\theta).$$



General Nested Filter

Algorithm

1. Initialisation

Draw $\theta_0^{(i)}, i = 1, \dots, N$, *i.i.d.* samples from $p(\theta)d\theta$.

2. Recursive step

2.1 For $i = 1, \dots, N$:

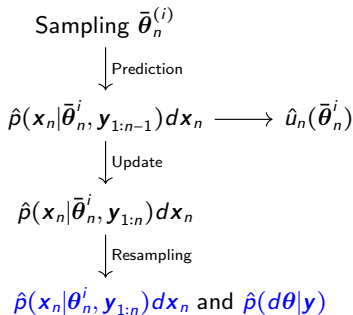
2.1.1 Draw $\bar{\theta}_n^{(i)}$ from $\kappa_N(d\theta|\theta_{n-1}^i)$.

2.1.2 Approximate $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$.

2.1.3 Use $p(\mathbf{x}_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})d\mathbf{x}_n$ to compute $\hat{u}_n(\bar{\theta}_n^i)$ and let $w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

2.2 Resample to obtain the set $\{\theta_{i=1}^N\}$ and the approximation

$$\hat{p}(d\theta|\mathbf{y}_{1:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}(d\theta).$$

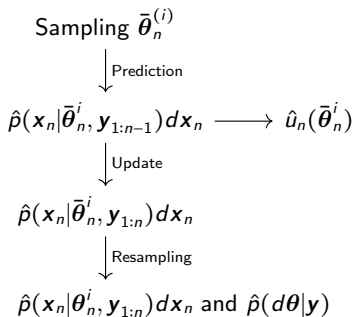


NHF via EKF

Let us assume $p(\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_0 | \bar{\mathbf{x}}_0, \bar{\mathbf{P}}_0)$,
 $\mathbf{v}_k \sim \mathcal{N}(\mathbf{v}_k | \mathbf{0}, \mathbf{Q})$ and $\mathbf{u}_k \sim \mathcal{N}(\mathbf{u}_k | \mathbf{0}, \mathbf{R})$.
 Functions $\bar{\mathbf{f}}$ in Eq. (2) and \mathbf{g} in Eq. (4) are
 nonlinear and differentiable, with $\mathbf{J}_{\bar{\mathbf{f}}, \mathbf{x}, \theta}$ and $\mathbf{J}_{\mathbf{g}, \mathbf{x}, \theta}$
 denoting their respective Jacobian matrices
 evaluated at the point \mathbf{x} and θ .

Algorithm

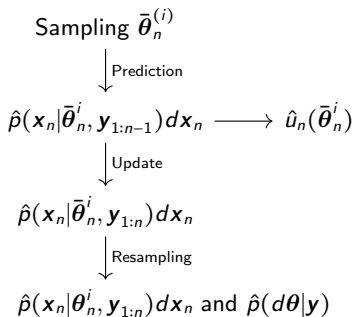
- Initialisation:* draw N i.i.d. particles
 $\theta_0^i \sim p(\theta) d\theta, i = 1, \dots, N$. Let $\bar{\mathbf{x}}_0^i = \bar{\mathbf{x}}_0$ and
 $\bar{\mathbf{P}}_0^i = \mathbf{P}_0$ for every i .
- Recursive step:* at time n , we have available
 $p(\theta | \mathbf{y}_{1:n-1}) d\theta \approx \frac{1}{N} \sum_{i=1}^N \delta_{\theta_{n-1}^i} (d\theta)$ and, for
 each $i = 1, \dots, N$, $p(\mathbf{x}_n | \theta_{n-1}^i, \mathbf{y}_{1:n-1}) d\mathbf{x}_n \approx$
 $\mathcal{N}(\mathbf{x}_n | \bar{\mathbf{x}}_{n-1}^i, \bar{\mathbf{P}}_{n-1}^i) d\mathbf{x}_n$.



NHF via EKF

Let us assume $p(\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_0 | \bar{\mathbf{x}}_0, \bar{\mathbf{P}}_0)$,
 $\mathbf{v}_k \sim \mathcal{N}(\mathbf{v}_k | \mathbf{0}, \mathbf{Q})$ and $\mathbf{u}_k \sim \mathcal{N}(\mathbf{u}_k | \mathbf{0}, \mathbf{R})$.
 Functions $\bar{\mathbf{f}}$ in Eq. (2) and \mathbf{g} in Eq. (4) are
 nonlinear and differentiable, with $\mathbf{J}_{\bar{\mathbf{f}}, \mathbf{x}, \theta}$ and $\mathbf{J}_{\mathbf{g}, \mathbf{x}, \theta}$
 denoting their respective Jacobian matrices
 evaluated at the point \mathbf{x} and θ .

Algorithm



- Initialisation:** draw N i.i.d. particles $\theta_0^i \sim p(\theta) d\theta, i = 1, \dots, N$. Let $\bar{\mathbf{x}}_0^i = \bar{\mathbf{x}}_0$ and $\bar{\mathbf{P}}_0^i = \mathbf{P}_0$ for every i .
- Recursive step:** at time n , we have available $p(\theta | \mathbf{y}_{1:n-1}) d\theta \approx \frac{1}{N} \sum_{i=1}^N \delta_{\bar{\theta}_{n-1}^i} (d\theta)$ and, for each $i = 1, \dots, N$, $p(\mathbf{x}_n | \theta_{n-1}^i, \mathbf{y}_{1:n-1}) d\mathbf{x}_{n-1} \approx \mathcal{N}(\mathbf{x}_{n-1} | \bar{\mathbf{x}}_{n-1}^i, \bar{\mathbf{P}}_{n-1}^i) d\mathbf{x}_n$.

(a) *Prediction:*

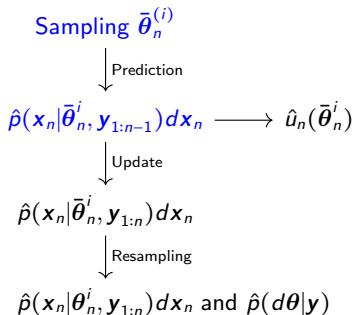
i Draw $\bar{\theta}_n^i \sim \kappa_N(d\theta|\theta_{n-1}^i)$, $i = 1, \dots, N$.

ii Let $\check{x}_0^i = \bar{x}_{n-1}^i$ and $\check{P}_0^i = \bar{P}_{n-1}^i$, $i = 1, \dots, N$. Then, for each i and $k = 1, \dots, T$ compute

$$\check{x}_k^i = \bar{f}_\sigma(\check{x}_{k-1}^i, \bar{\theta}_n^i, \mathbf{0}), \quad (9)$$

$$\check{P}_k^i = \mathbf{J}_{\bar{f}, \check{x}_{k-1}^i, \bar{\theta}_n^i} \check{P}_{k-1}^i \mathbf{J}_{\bar{f}, \check{x}_{k-1}^i, \bar{\theta}_n^i}^\top + \sigma^2 \mathbf{Q} \quad (10)$$

iii Set $p(x_n|\bar{\theta}_n^i, \mathbf{y}_{1:n-1})dx_n = \mathcal{N}(x_n|\hat{x}_n^i, \hat{P}_n^i)dx_n$ where $\hat{x}_n^i = \check{x}_T^i$ and $\hat{P}_n^i = \check{P}_T^i$.



(b) *Update:*

i For $i = 1, \dots, N$, compute

$$\mathbf{S}_n^i = \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i} \hat{\mathbf{P}}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}^\top + \sigma_o^2 \mathbf{R} \quad (11)$$

$$\mathbf{K}_n^i = \hat{\mathbf{P}}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}^\top (\mathbf{S}_n^i)^{-1} \quad (12)$$

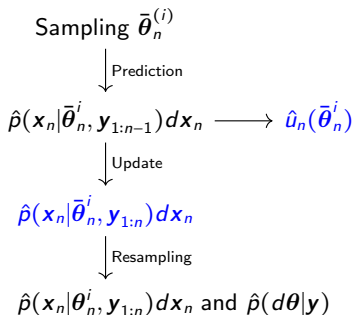
$$\check{\mathbf{x}}_n^i = \hat{\mathbf{x}}_n^i + \mathbf{K}_n^i (\mathbf{y}_n - \mathbf{g}(\hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i)) \quad (13)$$

$$\check{\mathbf{P}}_n^i = (\mathbf{I}_{d_x} - \mathbf{K}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}) \hat{\mathbf{P}}_n^i \quad (14)$$

ii Compute $\hat{\mathbf{u}}(\bar{\boldsymbol{\theta}}_n^i) = \mathcal{N}(\mathbf{y}_n | \mathbf{g}(\hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i), \mathbf{S}_n^i)$ and obtain the normalised weights.

iii Set the filter approximation

$$p(\mathbf{x}_n | \bar{\boldsymbol{\theta}}_n^i, \mathbf{y}_{1:n}) d\mathbf{x}_n = \mathcal{N}(\mathbf{x}_n | \check{\mathbf{x}}_n^i, \check{\mathbf{P}}_n^i) d\mathbf{x}_n.$$



(b) Update:

i For $i = 1, \dots, N$, compute

$$\mathbf{S}_n^i = \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i} \hat{\mathbf{P}}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}^\top + \sigma_o^2 \mathbf{R} \quad (11)$$

$$\mathbf{K}_n^i = \hat{\mathbf{P}}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}^\top (\mathbf{S}_n^i)^{-1} \quad (12)$$


$$\check{\mathbf{x}}_n^i = \hat{\mathbf{x}}_n^i + \mathbf{K}_n^i (\mathbf{y}_n - \mathbf{g}(\hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i)) \quad (13)$$

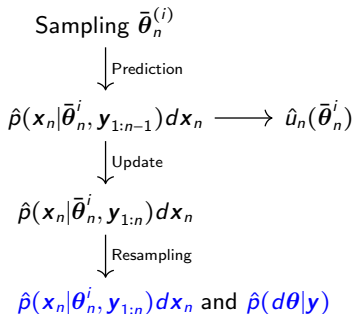
$$\check{\mathbf{P}}_n^i = (\mathbf{I}_{d_x} - \mathbf{K}_n^i \mathbf{J}_{\mathbf{g}, \hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i}) \hat{\mathbf{P}}_n^i \quad (14)$$

ii Compute $\hat{\mathbf{u}}(\bar{\boldsymbol{\theta}}_n^i) = \mathcal{N}(\mathbf{y}_n | \mathbf{g}(\hat{\mathbf{x}}_n^i, \bar{\boldsymbol{\theta}}_n^i), \mathbf{S}_n^i)$ and obtain the normalised weights.

iii Set the filter approximation

$$p(\mathbf{x}_n | \bar{\boldsymbol{\theta}}_n^i, \mathbf{y}_{1:n}) d\mathbf{x}_n = \mathcal{N}(\mathbf{x}_n | \check{\mathbf{x}}_n^i, \check{\mathbf{P}}_n^i) d\mathbf{x}_n.$$





(c) **Resampling**: draw indices j_1, \dots, j_N from the multinomial distribution with probabilities w_n^1, \dots, w_n^N , then set

$$\theta_n^i = \bar{\theta}_n^{j_i}, \quad \bar{\mathbf{x}}_n^i = \check{\mathbf{x}}_n^{j_i} \quad \text{and} \quad \bar{\mathbf{P}}_n^i = \check{\mathbf{P}}_n^{j_i} \quad (15)$$

for $i = 1, \dots, N$. Hence, $p(\mathbf{x}_n | \theta_n^i, \mathbf{y}_{1:n}) d\mathbf{x}_n = \mathcal{N}(\mathbf{x}_n | \bar{\mathbf{x}}_n^i, \bar{\mathbf{P}}_n^i) d\mathbf{x}_n$.

Computation of point estimates

- The posterior-mean estimators of θ and \mathbf{x} become

$$\hat{\theta}_n = \int \theta p(\theta | \mathbf{y}_{1:n}) d\theta \approx \frac{1}{N} \sum_{i=1}^N \theta_n^i = \hat{\theta}_n^N,$$

$$\mathbf{x}_n = \int \mathbf{x} p(\mathbf{x} | \mathbf{y}_{1:n}, \theta) d\mathbf{x}_n \approx \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{x}}_n^i = \mathbf{x}_n^N.$$

- The mean square error (MSE) of $\hat{\theta}_n$ is calculated as

$$\begin{aligned} \text{MSE}_n &= \int \|\theta - \hat{\theta}_n\|^2 p(\theta | \mathbf{y}_{1:n}) d\theta \\ &\approx \frac{1}{N} \sum_{i=1}^N \|\theta_n^i - \hat{\theta}_n^N\|^2 \end{aligned}$$

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

Preliminaries and notation

For the integral of a function $a(\boldsymbol{\theta})$ w.r.t. a measure α , hereafter we use the shorthand

$$(a, \alpha) := \int a(\boldsymbol{\theta})\alpha(d\boldsymbol{\theta}).$$

The sequence of posterior probability measures of the unknown parameters, $\mu_n(d\boldsymbol{\theta}) := p(\boldsymbol{\theta}|\mathbf{y}_{1:n})d\boldsymbol{\theta}$, $n \geq 1$, can be constructed recursively starting from a prior μ_0 as

$$\mu_n = u_n \star \mu_{n-1} \quad \text{where} \quad (f, u_n \star \alpha) = \frac{(fu_n, \alpha)}{(u_n, \alpha)}.$$

If, instead of the true likelihood u_n , we use another function $\bar{u}_n \neq u_n$ to update the posterior probability measure then we obtain the new sequence of measures

$$\bar{\mu}_0 = \mu_0, \quad \bar{\mu}_n = \bar{u}_n \star \bar{\mu}_{n-1}, \quad n = 1, 2, \dots$$

Convergence Theorem

A.1. The estimator $\hat{u}_n(\boldsymbol{\theta})$ is random and can be written as

$$\hat{u}_n(\boldsymbol{\theta}) = \bar{u}_n(\boldsymbol{\theta}) + m_n(\boldsymbol{\theta}),$$

where $m_n(\boldsymbol{\theta})$ is a zero-mean random variable with finite variance. Furthermore, the mean $\bar{u}_n(\boldsymbol{\theta}) = \mathbb{E}[\hat{u}_n(\boldsymbol{\theta})]$ has the form

$$\bar{u}_n(\boldsymbol{\theta}) = u_n(\boldsymbol{\theta}) + b_n(\boldsymbol{\theta}),$$

where $b_n(\boldsymbol{\theta})$ is a deterministic and bounded bias function.

Theorem 1

Let the sequence of observations $y_{1:n_o}$ be arbitrary but fixed, with $n_o < \infty$, and choose an arbitrary function $h \in B(D)$. Let $\mu_n^N = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_i^n}$ be the random probability measure in the parameter space generated by the nested filter. If A.1 holds and under regularity conditions, then

$$\|(h, \mu_n^N) - (h, \bar{\mu}_n)\|_p \leq \frac{c_n \|h\|_\infty}{\sqrt{N}}, \quad \text{for } n = 0, 1, \dots, n_o,$$

where $\{c_n\}_{0 \leq n \leq n_o}$ is a sequence of constants independent of N . \square

Convergence Theorem

A.1. The estimator $\hat{u}_n(\theta)$ is random and can be written as

$$\hat{u}_n(\theta) = \bar{u}_n(\theta) + m_n(\theta),$$

where $m_n(\theta)$ is a zero-mean random variable with finite variance. Furthermore, the mean $\bar{u}_n(\theta) = E[\hat{u}_n(\theta)]$ has the form

$$\bar{u}_n(\theta) = u_n(\theta) + b_n(\theta),$$

where $b_n(\theta)$ is a deterministic and bounded bias function.

Theorem 1

Let the sequence of observations $y_{1:n_o}$ be arbitrary but fixed, with $n_o < \infty$, and choose an arbitrary function $h \in B(D)$. Let $\mu_n^N = \frac{1}{N} \sum_{i=1}^N \delta_{\theta_n^i}$ be the random probability measure in the parameter space generated by the nested filter. If A.1 holds and under regularity conditions, then

$$\|(h, \mu_n^N) - (h, \bar{\mu}_n)\|_p \leq \frac{c_n \|h\|_\infty}{\sqrt{N}}, \quad \text{for } n = 0, 1, \dots, n_o,$$

where $\{c_n\}_{0 \leq n \leq n_o}$ is a sequence of constants independent of N . \square

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

A Stochastic Lorenz 96 Model

- The model consists of two sets of dynamic variables, $\mathbf{x}(t)$ and $\mathbf{z}(t)$, that displays some key features of atmosphere dynamics (including chaotic behaviour) in a relatively simple model of arbitrary dimension. The system of differential equations takes the form

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \mathbf{f}_1(\mathbf{x}(t), \mathbf{z}(t), \boldsymbol{\alpha})dt + dW \\ \dot{\mathbf{z}}(t) &= \mathbf{f}_2(\mathbf{x}(t), \mathbf{z}(t), \boldsymbol{\alpha})dt + d\bar{W}\end{aligned}\quad (16)$$

- Let us assume there are d_x slow variables and L fast variables per slow variable. The maps, \mathbf{f}_1 and \mathbf{f}_2 functions, can be written as

$$\begin{aligned}\mathbf{f}_1 &= [f_{11}, \dots, f_{1d_x}]^T, \\ \mathbf{f}_2 &= [f_{21}, \dots, f_{2L}]^T,\end{aligned}$$

where

$$\begin{aligned}f_{1j}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -x_{j-1}(x_{j-2} - x_{j+1}) - x_j + F - \frac{HC}{B} \sum_{l=(j-1)L}^{Lj-1} z_l, \\ f_{2l}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -CBz_{l+1}(z_{l+2} - z_{l-1}) - Cz_l + \frac{CF}{B} + \frac{HC}{B} x_{\lfloor \frac{l-1}{L} \rfloor}.\end{aligned}$$

A Stochastic Lorenz 96 Model

- The model consists of two sets of dynamic variables, $\mathbf{x}(t)$ and $\mathbf{z}(t)$, that displays some key features of atmosphere dynamics (including chaotic behaviour) in a relatively simple model of arbitrary dimension. The system of differential equations takes the form

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \mathbf{f}_1(\mathbf{x}(t), \mathbf{z}(t), \boldsymbol{\alpha})dt + dW \\ \dot{\mathbf{z}}(t) &= \mathbf{f}_2(\mathbf{x}(t), \mathbf{z}(t), \boldsymbol{\alpha})dt + d\bar{W}\end{aligned}\quad (16)$$

- Let us assume there are d_x slow variables and L fast variables per slow variable. The maps, \mathbf{f}_1 and \mathbf{f}_2 functions, can be written as

$$\begin{aligned}\mathbf{f}_1 &= [f_{11}, \dots, f_{1d_x}]^T, \\ \mathbf{f}_2 &= [f_{21}, \dots, f_{2L}]^T,\end{aligned}$$

where

$$\begin{aligned}f_{1j}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -x_{j-1}(x_{j-2} - x_{j+1}) - x_j + F - \frac{HC}{B} \sum_{l=(j-1)L}^{Lj-1} z_l, \\ f_{2l}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -CBz_{l+1}(z_{l+2} - z_{l-1}) - Cz_l + \frac{CF}{B} + \frac{HC}{B} x_{\lfloor \frac{l-1}{L} \rfloor}.\end{aligned}$$

A Stochastic Lorenz 96 Model

- Applying the a discretization method we obtain a discrete-time version of the two-scale Lorenz 96 model

$$\bar{\mathbf{x}}_k = \bar{\mathbf{x}}_{k-1} + h\bar{\mathbf{f}}_{1,\sigma}(\bar{\mathbf{x}}_{k-1}, \bar{\mathbf{z}}_{k-1}, \boldsymbol{\alpha}, \mathbf{v}_k), \quad (17)$$

$$\bar{\mathbf{z}}_k = \bar{\mathbf{z}}_{k-1} + h\bar{\mathbf{f}}_{2,\bar{\sigma}}(\bar{\mathbf{x}}_{k-1}, \bar{\mathbf{z}}_{k-1}, \boldsymbol{\alpha}, \bar{\mathbf{v}}_k) \quad (18)$$

- We assume that the observations are linear but can only be collected from this system once every T discrete time steps and only one out of K slow variables can be observed, having the form

$$\mathbf{y}_n = \begin{bmatrix} x_{K,nT} \\ x_{2K,nT} \\ \vdots \\ x_{d_y K,nT} \end{bmatrix} + \mathbf{u}_n, \quad (19)$$

A Stochastic Lorenz 96 Model

- Applying the a discretization method we obtain a discrete-time version of the two-scale Lorenz 96 model

$$\bar{\mathbf{x}}_k = \bar{\mathbf{x}}_{k-1} + h\bar{\mathbf{f}}_{1,\sigma}(\bar{\mathbf{x}}_{k-1}, \bar{\mathbf{z}}_{k-1}, \boldsymbol{\alpha}, \mathbf{v}_k), \quad (17)$$

$$\bar{\mathbf{z}}_k = \bar{\mathbf{z}}_{k-1} + h\bar{\mathbf{f}}_{2,\bar{\sigma}}(\bar{\mathbf{x}}_{k-1}, \bar{\mathbf{z}}_{k-1}, \boldsymbol{\alpha}, \bar{\mathbf{v}}_k) \quad (18)$$

- We assume that the observations are linear but can only be collected from this system once every T discrete time steps and only one out of K slow variables can be observed, having the form

$$\mathbf{y}_n = \begin{bmatrix} x_{K,nT} \\ x_{2K,nT} \\ \vdots \\ x_{d_y K,nT} \end{bmatrix} + \mathbf{u}_n, \quad (19)$$

A Stochastic Lorenz 96 Model

- We generate both ground-truth values for the slow variables $\{\mathbf{x}_n\}_{n \geq 0}$ and synthetic observations, $\{\mathbf{y}_n\}_{n \geq 1}$. In the algorithm, as a forecast model, we use the differential equation

$$\dot{x}_j = f_j(\mathbf{x}, \boldsymbol{\theta}) = -x_{j-1}(x_{j-2} - x_{j+1}) - x_j + F - \ell(x_j, \mathbf{a}) \quad (20)$$

where

- $\mathbf{a} = [a_1, a_2]^\top$ is a (constant) parameter vector,
- $\boldsymbol{\theta} = [F, \mathbf{a}^\top]^\top$ contains all the parameters and
- function $\ell(x_j, \mathbf{a})$ is an ansatz, a polynomial in x_j of degree 2, for the coupling term $\frac{HC}{B} \sum_{l=(j-1)L}^{Lj-1} \bar{z}_l$.
- Applying a discretization method we obtain

$$\mathbf{x}_k = \mathbf{x}_{k-1} + h\bar{\mathbf{f}}_\sigma(\mathbf{x}_{k-1}, \boldsymbol{\theta}, \mathbf{v}_k) \quad (21)$$

Simulation setup

$$\begin{aligned}
 f_{1j}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -x_{j-1}(x_{j-2} - x_{j+1}) - x_j + F - \frac{HC}{B} \sum_{l=(j-1)L}^{Lj-1} z_l, \\
 f_{2l}(\mathbf{x}, \mathbf{z}, \boldsymbol{\alpha}) &= -CBz_{l+1}(z_{l+2} - z_{l-1}) - Cz_l + \frac{CF}{B} + \frac{HC}{B} x_{\lfloor \frac{l-1}{L} \rfloor}.
 \end{aligned} \quad (22)$$

Integration step	$h = 10^{-3}$
Model parameters	$F = 8, H = 0.75, C = 10$ and $B = 15$
Fast variables	$L = 10$
Observed variables	$K = 2$
Noise scaling factors	$\sigma = \frac{h}{4} = 0.25 \times 10^{-3}$ and $\sigma_o = 4$

- The accuracy of the estimation is assessed in terms of empirical MSE per dimension averaged over several independent simulation runs.

$$\text{MSE}_k = \frac{1}{d_x} \|\check{\mathbf{x}}_k - \tilde{\mathbf{x}}_k\|^2. \quad (23)$$

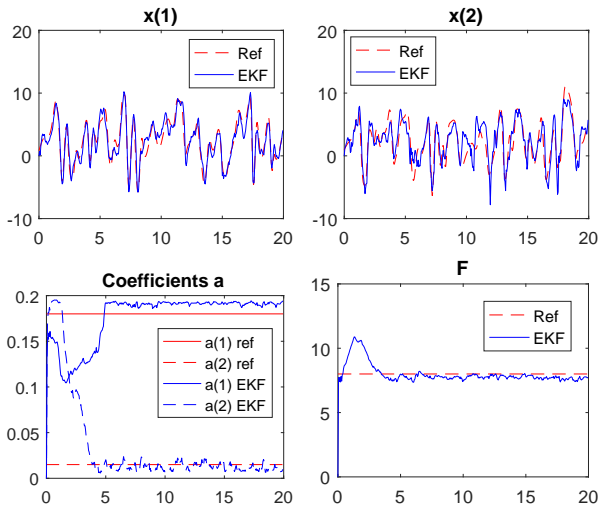
NPF vs NHF's

- Dimension $d_x = 50$ and gap between observations of $hT = 0.05$ continuous-time units.
- $N = M = 1400$ particles in both layers of NPF.
- $N = 200$ particles for the first layer of NHFs and $M = 50$ for the EnKF's run.

Algorithm	Running time	MSE
NPF	6.85 hours	1.872
NHF + EKF	1.196 minutes	1.653
NHF + EnKF	11.674 minutes	0.472

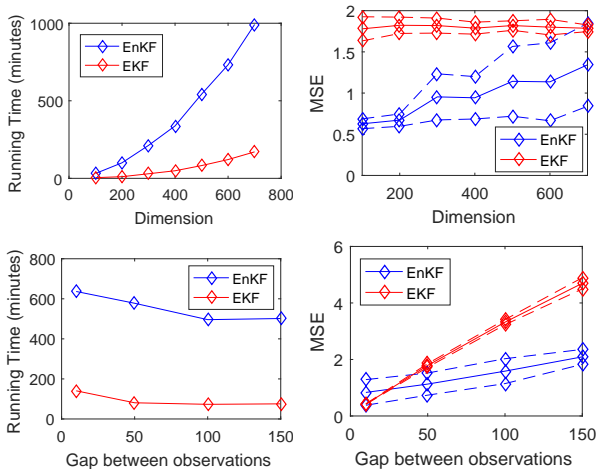
Table: Running times and average MSE (over time and state dimensions) for the NPF and two NHFs, based on the EKF and the EnKF, respectively.

4000-dimensional Experiment



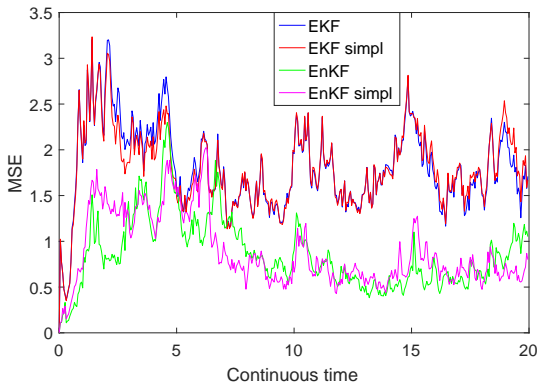
Sequences of state estimates in x_1 and x_2 and estimates of the parameters $a = [a_1, a_2]^T$ and F over time in a 4,000-dimensional Lorenz 96 model. Variable x_1 is observed (in Gaussian noise), while x_2 is unobserved. The reference values are represented in red lines.

NHF-EKF vs NHF-EnKF



Comparison of the NHF-EKF (red lines) and NHF-EnKF (blue lines) in terms of their running time and their MSE as the state dimension d_x increase ($T = 50$ discrete time steps) and as the gap between observations T increases ($d_x = 500$).

Simplifications Performance



Comparison of the original scheme of NHF using both Gaussian filters (EKF and EnKF) and the inverse simplifications.

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

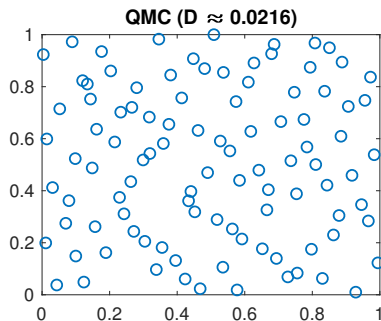

QMC

- We are going to study an alternative to MC, that works similarly but in a **deterministic way**.
- The main difference with MC is the use of a **low-discrepancy** point set (sub-random or quasi-random) to generate the samples.
- Discrepancy is defined as

$$D(\mathbf{u}^{(1:N)}, \mathcal{A}) = \sup_{A \in \mathcal{A}} \left| \frac{1}{N} \sum_{n=1}^N \mathbf{1}(\mathbf{u}^n \in [a, b]) - \lambda_s(A) \right|$$

where we have N vectors $\mathbf{u}^n \in [0, 1)^d$, $\lambda_s(A)$ is the volume of A and \mathcal{A} is the family of measurable subsets of $[0, 1)^d$.

QMC



QMC vs MC: $N = 100$ points sampled independently and uniformly in $[0, 1]^2$ (left) and QMC sequence in $[0, 1]^2$ of the same length (right).

Discrepancy preserving map


- $\Phi : D_\theta \rightarrow [0, 1)^d$ is a mapping that is discrepancy preserving, where D_θ is the support of θ .
- Then, $\Phi(\theta) = (\Phi_1(\theta_1), \dots, \Phi_n(\theta_n))$, where Φ_i 's are continuous and strictly monotone, e.g.,

$$\Phi_i(\theta_i) = \left[1 + \exp\left(-\frac{\theta_i - \underline{\theta}_i}{\bar{\theta}_i - \underline{\theta}_i}\right) \right]^{-1}, \quad i = 1, \dots, d$$

where $\underline{\theta}_i = \mu_{\theta_{1:N}} - 2\sigma_{\theta_{1:N}}$ and $\bar{\theta}_i = \mu_{\theta_{1:N}} + 2\sigma_{\theta_{1:N}}$.

Hilbert curve

- The Hilbert curve is a continuous fractal map $H : [0, 1) \rightarrow [0, 1)^d$.
- It admits a pseudo-inverse $h : [0, 1)^d \rightarrow [0, 1)$, which is used in SQMC to sort the samples $\bar{\theta}_i$.



Hilbert curves of order $m = 4, 8, 16$.

Sequential QMC

Algorithm

1. Initialisation

- 1.1 Generate a QMC point set $\mathbf{v}_0^{(i)}$ in $(0, 1]^d$ and compute $\bar{\theta}_0^{(i)}$ from $\kappa_N(d\theta, \mathbf{v}_0^{(i)})$.
- 1.2 Compute the normalised weights $w_0^i \propto \hat{u}_0(\bar{\theta}_0^i)$.

2. Recursive step

- 2.1 Generate a QMC point set $\mathbf{V}_n^{(i)} = [r_n^{(i)}, \mathbf{v}_n^{(i)}]$ in $(0, 1]^{d+1}$.
- 2.2 Hilbert sort: find permutation σ_{t-1} such that $h \circ \phi(\theta_n^{\sigma_{t-1}(1)}) \leq \dots \leq h \circ \phi(\theta_n^{\sigma_{t-1}(N)})$ and reorder the weights $w_n^{\sigma_{t-1}(i)}$.
- 2.3 Find permutation τ such that $r_t^{\tau(1)} \leq \dots \leq r_t^{\tau(N)}$, and obtain the set $\{\theta_{i=1}^N\}$ applying the inversion method. Compute $\bar{\theta}_{n+1}^{(i)}$ from $\kappa_N(d\theta | \theta_n^i, \mathbf{v}_n^i)$.
- 2.4 Compute $\hat{u}_n(\bar{\theta}_n^i)$ and let $w_n^i \propto \hat{u}_n(\bar{\theta}_n^i)$ be the normalised weight of $\bar{\theta}_n^i$.

Lorenz 63 model

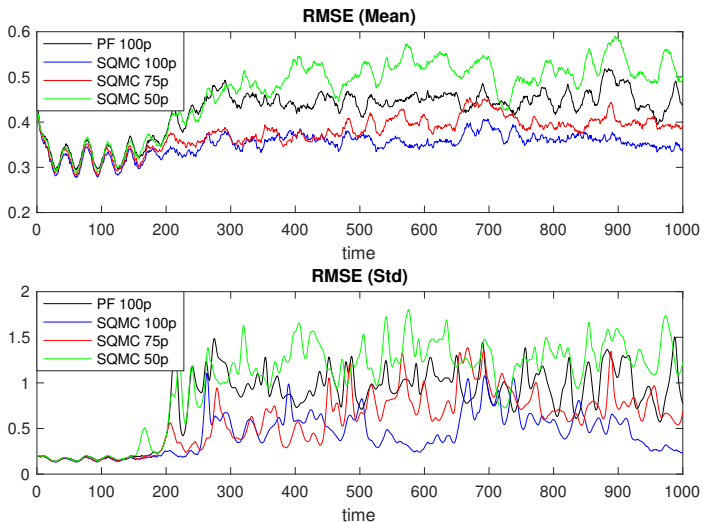
- In this model the state-space consists in three variables (x , y and z), each of them with its respective SDE.

$$f_x(x, y, z, S) = -S(x - y),$$

$$f_y(x, y, z, R) = Rx - y - xz,$$

$$f_z(x, y, z, B) = xy - Bz,$$

- In order to do some experiments, they are discretized and the parameters are considered known ($S = 10$, $R = 28$ and $Q = \frac{8}{3}$).



SQMC vs SMC. Results obtained over 1000 simulation runs of a Lorenz 63 model.

Index

Introduction

State of the art and contributions

Nested Hybrid Filters

Dynamical Model

Nested Filters

Convergence Analysis

Convergence theorem and analysis of the divergence between μ and $\bar{\mu}$

Numerical results

Lorenz 96 Model

Simulation setup

Results

SQMC

Getting into some details

Sequential Quasi-Monte Carlo

Conclusions

Concluding remarks

Future work

Conclusions

- We have introduced a **recursive** methodology to estimate the static parameters and the dynamic variables.
- The use of Gaussian filters is investigated as they admit **fast implementations** that can be well suited to **high dimensional** systems and two of them are simulated.
- **Simplifications** in the NHF-EKF and NHF-EnKF allowed the implementation of **high dimensional systems**.
- We have proved, under very general assumptions, that **the proposed method converges** (with optimal Monte Carlo rates) to a possibly biased version of the posterior distribution of the parameters.

Future work

- Apply SQMC and other filters in the first layer of the algorithm in order to improve its performance.
- Use this algorithms with [real data](#) thanks to our collaboration with MeteoGalicia.