Backward sequential Monte Carlo for marginal smoothing

Joel Kronander*, Thomas B. Schön† and Johan Dahlin‡

LINKÖPINGS UNIVERSITET · UPPSALA UNIVERSITET

Summary

- We propose a new forward-backward particle smoother with linear computational complexity.
- The new smoother is based on running a Sequential Monte Carlo sampler backwards in time after an initial forward filtering pass.
- We show through simulation studies that the new smoother can outperform existing smoothers when targeting the marginal smoothing densities.

Forward-Backward Particle Smoothers

Consider a general **nonlinear state space model (SSM)**

$$x_{t+1}|x_t \sim f(x_{t+1}|x_t)$$
$$y_t|x_t \sim g(y_t|x_t)$$

We are interested in inferring the **marginal smoothing distribution**, $p(x_t|y_{1:T})$. Forward-backward smoothers are based on sampling M backward paths after an initial forward particle filter, providing approximations of the filtering distributions $p(x_t|y_{1:t})$.

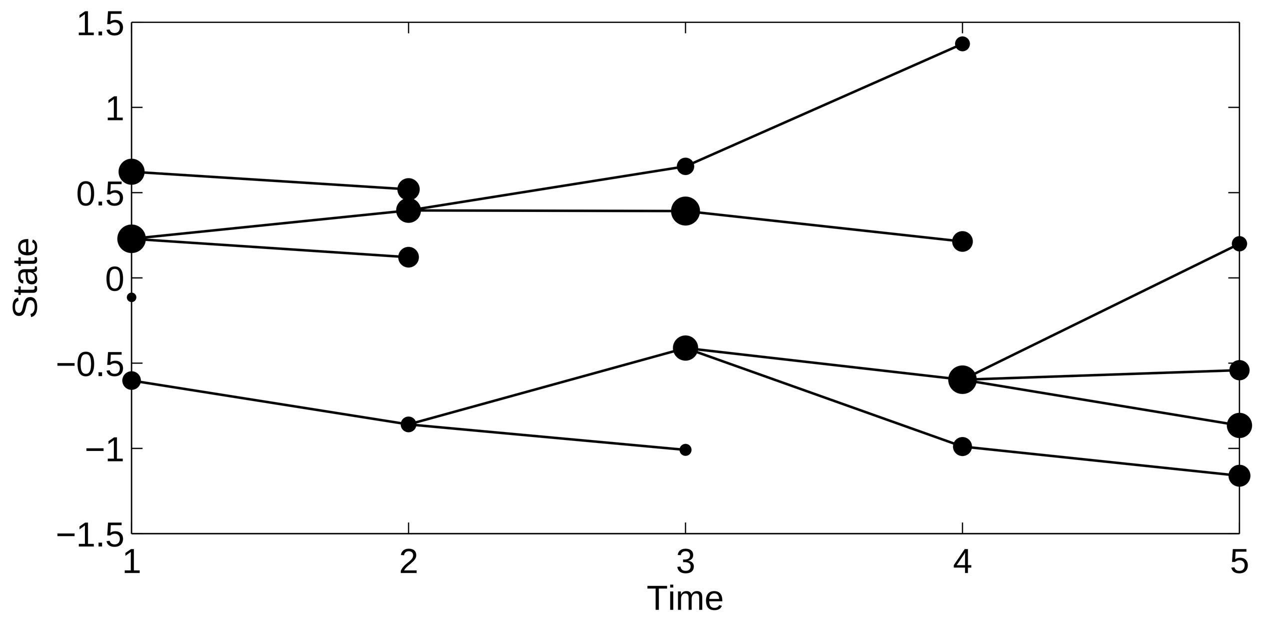

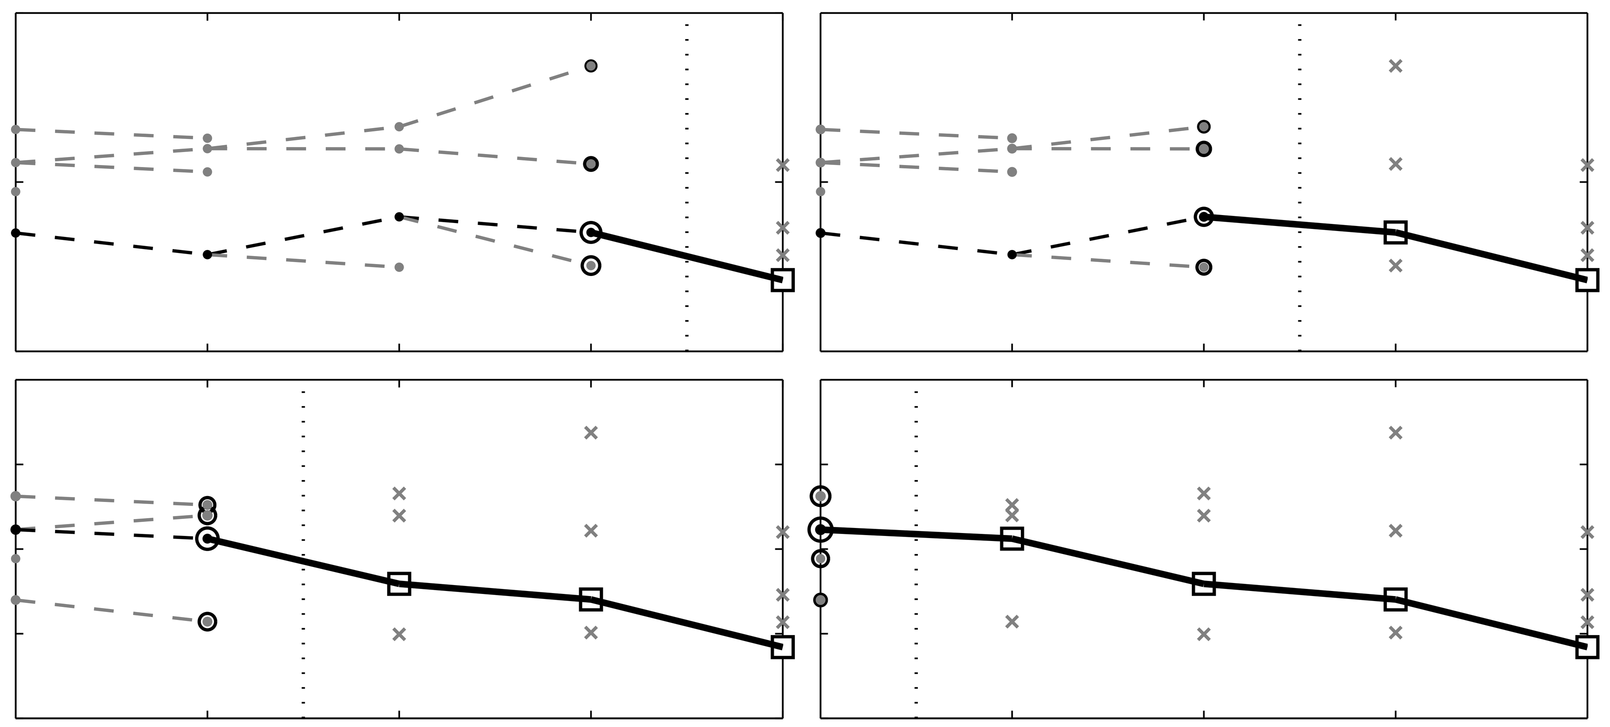

Main idea

Improve previous forward-backward smoothers that samples independent backward paths by using a Sequential Monte Carlo approach to exploit information gathered in other backward paths.

SMC smoother

We can express marginal smoothing distribution at time t as

$$p(x_t|y_{1:T}) = \int \frac{p(x_t|y_{1:t})g(y_{t+1}|x_{t+1})f(x_{t+1}|x_t)p(x_{t+1}|y_{1:T})}{p(x_{t+1}|y_{1:t+1})p(y_{t+1}|y_{1:t})} dx_{t+1} \quad (1)$$

By using the empirical approximation of the forward filtering density at time t and the marginal smoothing densities at time $t+1$ and assuming that the backward particle system at time $t+1$ is in the support of the forward filter particle system we can approximate (1) by

$$\widehat{p}(x_t|y_{1:T}) \propto \sum_{i=1}^{N}\sum_{j=1}^{M} w_t^i \widetilde{w}_{t+1}^j \frac{g(y_{t+1}|\widetilde{x}_{t+1}^j)f(\widetilde{x}_{t+1}^j|x_t^i)}{p(\widetilde{x}_{t+1}^j|y_{1:t+1})}\delta_{x_t^i}(x_t) \quad (2)$$

$$\widehat{p}(\widetilde{x}_{t+1}^j|y_{1:t+1}) = \begin{cases} w_{t+1}^i & \text{if } \widetilde{x}_{t+1}^j = x_{t+1}^i \\ 0 & \text{if } \widetilde{x}_{t+1}^j \neq x_{t+1}^i \end{cases} \quad (3)$$

This allow us to target (2) using importance sampling similar in spirit to the auxiliary particle filter.

Backward SMC smoother

- Sample M backward particles, $\widetilde{x}_T^m$ from the forward particles at time T, $\{x_T^i, w_T^i\}_{i=1:N}$
For $t = T-1, \ldots, 1$
    For $j = 1, \ldots, M$
        - Sample $a_t^j \sim Cat\left(\{w_t^i\}_{i=1}^N\right)$
        - Sample $b_t^j \sim Cat\left(\left\{\frac{\widetilde{w}_{t+1}^k g(y_{t+1}|\widetilde{x}_{t+1}^{k_t})}{\widehat{p}(\widetilde{x}_{t+1}^{k_t}|y_{1:t+1})}\right\}_{k=1}^M\right)$
        - Set $\widetilde{x}_t^j = x_t^{a_t^j}$
        - Set $\widetilde{w}_t^j = f(\widetilde{x}_{t+1}^{b_t^j}|\widetilde{x}_t^j)$
    End For
    - Set $\widetilde{w}_t^j = \frac{\widetilde{w}_t^j}{\sum_{j=1}^M \widetilde{w}_t^j}$ for $j = 1, \ldots, M$
End For

Numerical Illustration

Consider a generic ten dimensional linear Gaussian SSM

$$x_{t+1}|x_t \sim \mathcal{N}(Ax_t, Q) \quad (4a)$$
$$y_t|x_t \sim \mathcal{N}(Cx_t, R) \quad (4b)$$

For this model 50 independent realizations of A and C were generated using the MATLAB command `drss(10,10,0)`, and Q and R were set to the identity matrix.

	FFBSi	RS-FFBSi	MH-FFBSi	SMC Smoother
MSE	0.66	0.66	0.66	0.66
Runtime	2.02	1.75	0.38	**0.10**

Average mean squared error compared to the Rauch-Tung-Striebel smoother and runtime for linear Gaussian model (4), using $N = 200$ forward filtering particles and $M = 100$ backward particles.

Altough the new smoother performs well in practice with reasonable particles counts, we have found that it exhibits a small non-vanishing bias. In our current efforts we are investigating methods to quantify and eliminate this bias.

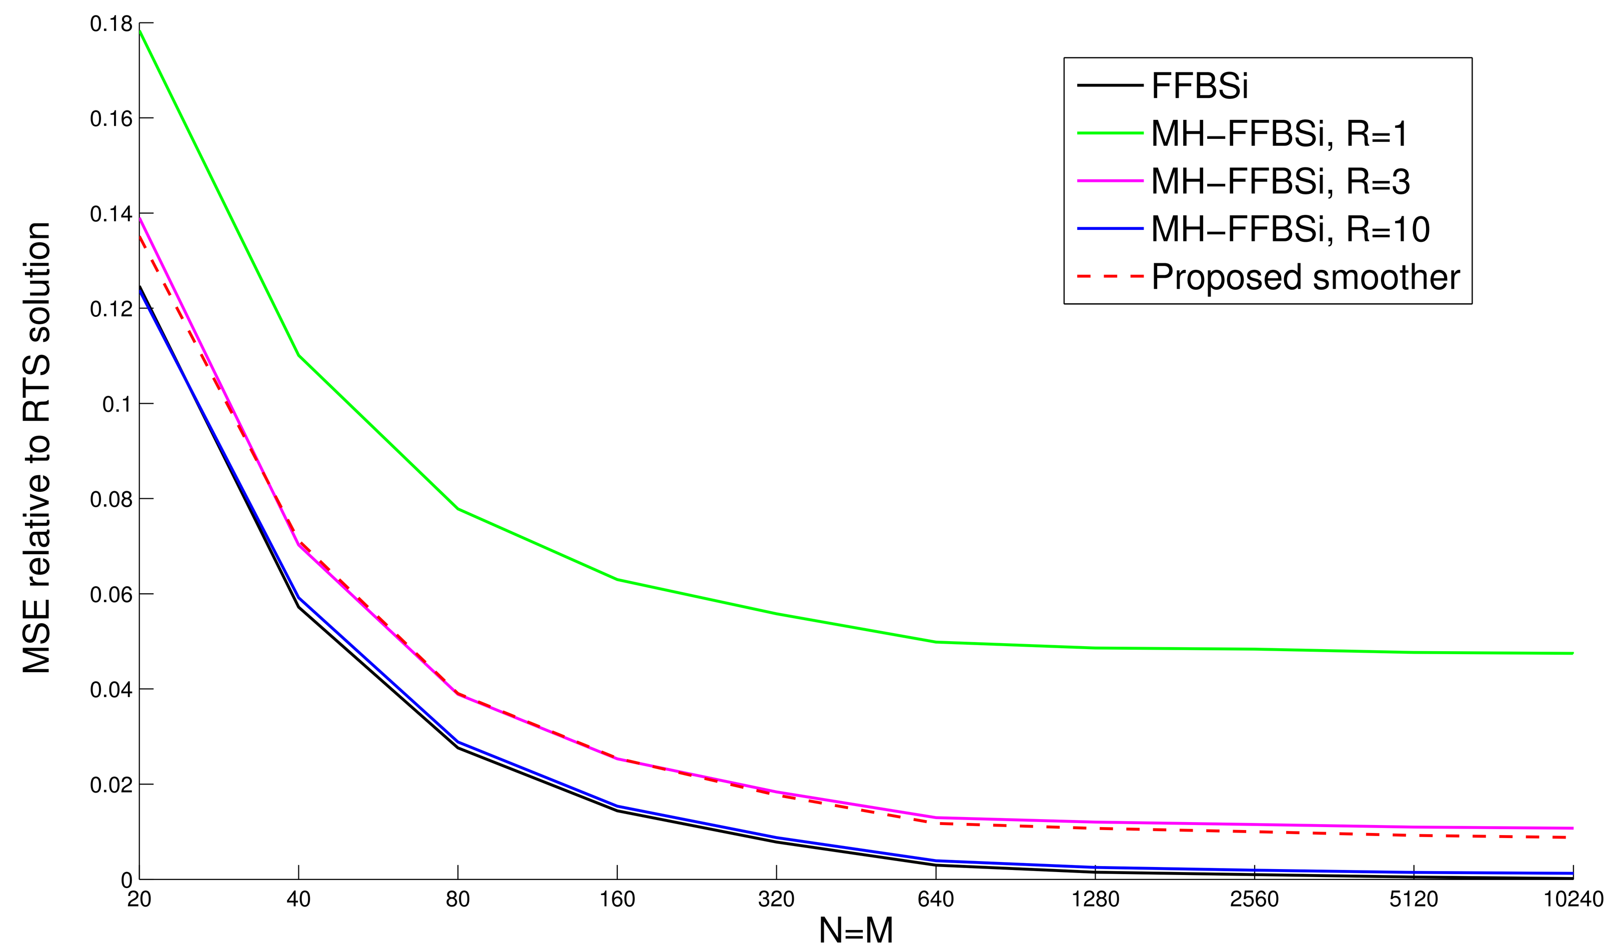

Non-vanishing bias for a one dimensional linear Gaussian SSM.

More information and source code

http://vcl.itn.liu.se/publications/2014/KBD14/

*joel.kronander@liu.se, †thomas.schon@it.uu.se and ‡johan.dahlin@liu.se