## Learning dynamical systems using SMC

Thomas Schön, Uppsala University

Max Planck Institute for Intelligent Systems
Tübingen, Germany
May 29, 2018.

## What we do in the team

## We automate the extraction of knowledge and understanding from data.

Both basic research and applied research (with companies).


Create probabilistic models of dynamical systems and their surroundings.

Develop methods to learn models from data.
The models can then be used by machines (or humans) to understand or take decisions about what will happen next.

## Nonlinear state space model (SSM)

The state space model (SSM) is a Markov chain that makes use of a latent variable representation to describe dynamical phenomena.

It consists of two stochastic processes:

1. unobserved (state) process $\left\{x_{t}\right\}_{t \geq 0}$ modelling the dynamics,
2. observed process $\left\{y_{t}\right\}_{t \geq 1}$ modelling the measurements and their relationship to the unobserved state process.

$$
\begin{aligned}
& x_{t}=f\left(x_{t-1}, \theta\right)+v_{t}, \\
& y_{t}=g\left(x_{t}, \theta\right)+e_{t},
\end{aligned}
$$

where $\theta \in \mathbb{R}^{n_{\theta}}$ denotes static model parameters.
The SSM offers a practical representation not only for modelling, but also for reasoning and inference.

## Ex) "what are $x_{t}, \theta$ and $y_{t}$ "?

Aim (motion capture): Compute $x_{t}$ (position and orientation of the different body segments) of a person ( $\theta$ describes the body shape) moving around indoors using measurements $y_{t}$ (accelerometers, gyroscopes and ultrawideband).


## Show movie!

[^0]
## Three different representations of the SSM

Three alternative representations, using

1. graphical models,
2. probability distributions or
3. probabilistic programs.
4. Representing the SSM using a graphical model:


## Representations using distributions or probabilistic programs

2. Representation using probability distributions

$$
\begin{aligned}
x_{t} \mid\left(x_{t-1}, \theta\right) & \sim p\left(x_{t} \mid x_{t-1}, \theta\right) \\
y_{t} \mid\left(x_{t}, \theta\right) & \sim p\left(y_{t} \mid x_{t}, \theta\right) \\
x_{0} & \sim p\left(x_{0} \mid \theta\right)
\end{aligned}
$$

3. Representing the SSM using a probabilistic program

$$
\begin{aligned}
\mathrm{x}[1] \sim \operatorname{Gaussian}(0.0,1.0) ; & p\left(x_{1}\right) \\
\mathrm{y}[1] \sim \operatorname{Gaussian}(\mathrm{x}[1], 1.0) ; & p\left(y_{1} \mid x_{1}\right) \\
\text { for }(\mathrm{t} \text { in 2..T })\{ & \\
\mathrm{x}[\mathrm{t}] \sim \operatorname{Gaussian}\left(\mathrm{a}^{*} \mathrm{x}[\mathrm{t}-1], 1.0\right) ; & p\left(x_{t} \mid x_{t-1}\right) \\
\mathrm{y}[\mathrm{t}] & \sim \operatorname{Gaussian}(\mathrm{x}[\mathrm{t}], 1.0) ;
\end{aligned}
$$

$$
\}
$$

A probabilistic program encodes a probabilistic model (here an LG-SSM) according to the semantics of a particular probabilistic programming language (here Birch).

## SSM - full probabilistic model

The full probabilistic model is given by

$$
p\left(x_{0: T}, \theta, y_{1: T}\right)=\underbrace{p\left(y_{1: T} \mid x_{0: T}, \theta\right)}_{\text {data distribution }} \underbrace{p\left(x_{0: T}, \theta\right)}_{\text {prior }}
$$

Distribution describing a parametric nonlinear SSM

$$
p\left(x_{0: T}, \theta, y_{1: T}\right)=\underbrace{\prod_{t=1}^{T} \underbrace{p\left(y_{t} \mid x_{t}, \theta\right)}_{\text {observation }}}_{\text {data distribution }} \underbrace{\prod_{t=1}^{T} \underbrace{p\left(x_{t} \mid x_{t-1}, \theta\right)}_{\text {dynamics }} \underbrace{p\left(x_{0} \mid \theta\right)}_{\text {state }} \underbrace{p(\theta)}_{\text {param. }}}_{\text {prior }}
$$

Model $=$ probability distribution!

## Learning the states and the parameters

Based on our generative model, compute the posterior distribution

$$
p\left(x_{0: T}, \theta \mid y_{1: T}\right)=\underbrace{p\left(x_{0: T} \mid \theta, y_{1: T}\right)}_{\text {state inf. }} \underbrace{p\left(\theta \mid y_{1: T}\right)}_{\text {param. inf. }}
$$

Bayesian formulation - model the unknown parameters as a random variable $\theta \sim p(\theta)$ and compute

$$
p\left(\theta \mid y_{1: T}\right)=\frac{p\left(y_{1: T} \mid \theta\right) p(\theta)}{p\left(y_{1: T}\right)}
$$

Maximum likelihood formulation - model the unknown parameters as a deterministic variable and solve

$$
\widehat{\theta}=\underset{\theta \in \Theta}{\arg \max } p\left(y_{1: T} \mid \theta\right) .
$$

## Central object - the likelihood

The likelihood is computed by marginalizing

$$
p\left(x_{0: T}, y_{1: T} \mid \theta\right)=p\left(x_{0} \mid \theta\right) \prod_{t=1}^{T} p\left(y_{t} \mid x_{t}, \theta\right) \prod_{t=1}^{T} p\left(x_{t} \mid x_{t-1}, \theta\right)
$$

w.r.t the state sequence $x_{0: T}$,

$$
p\left(y_{1: T} \mid \theta\right)=\int p\left(x_{0: T}, y_{1: T} \mid \theta\right) \mathrm{d} x_{0: T}
$$

(We are averaging $p\left(x_{0: T}, y_{1: T} \mid \theta\right)$ over all possible state sequences.)
Equivalently we have

$$
p\left(y_{1: T} \mid \theta\right)=\prod_{t=1}^{T} p\left(y_{t} \mid y_{1: t-1}, \theta\right)=\prod_{t=1}^{T} \int p\left(y_{t} \mid x_{t}, \theta\right) \underbrace{p\left(x_{t} \mid y_{1: t-1}, \theta\right)}_{\text {key challenge }} \mathrm{d} x_{t} .
$$

[^1]
## State inference - nonlinear filtering problem

The nonlinear filtering problem involves the measurement update

$$
p\left(x_{t} \mid y_{1: t}\right)=\frac{\overbrace{p\left(y_{t} \mid x_{t}\right)}^{\text {measurement }} \overbrace{p\left(x_{t} \mid y_{1: t-1}\right)}^{\text {prediction pdf }}}{p\left(y_{t} \mid y_{1: t-1}\right)},
$$

and the time update

$$
p\left(x_{t} \mid y_{1: t-1}\right)=\int \underbrace{p\left(x_{t} \mid x_{t-1}\right)}_{\text {dynamics }} \underbrace{p\left(x_{t-1} \mid y_{1: t-1}\right)}_{\text {filtering pdf }} \mathrm{d} x_{t-1}
$$

## Outline

Aim: 1. Give a (hopefully) intuitive explanation of sequential Monte Carlo (SMC) for probabilistic modelling of dynamical systems.
2. Derive a new stochastic optimization method that can for example be used to learn unknown parameters in nonlinear SSMs.

## 1. Probabilistic modelling of dynamical systems

2. Sequential Monte Carlo (SMC)
3. Stochastic optimization
4. Some ongoing research snapshots (if there is time)

## Sequential Monte Carlo

The need for computational methods, such as SMC, is tightly coupled to the intractability of the integrals on the previous slides.

SMC provide approximate solutions to integration problems where there is a sequential structure present.

The particle filter approximates $p\left(x_{t} \mid y_{1: t}\right)$ for

$$
\begin{aligned}
& x_{t}=f\left(x_{t-1}\right)+v_{t}, \\
& y_{t}=g\left(x_{t}\right)+e_{t},
\end{aligned}
$$

by maintaining an empirical distribution made up of $N$ samples (particles) $\left\{x_{t}^{i}\right\}_{i=1}^{N}$ and the corresponding weights $\left\{w_{t}^{i}\right\}_{i=1}^{N}$

$$
\underbrace{\widehat{p}\left(x_{t} \mid y_{1: t}\right)}_{\widehat{\pi}\left(x_{t}\right)}=\sum_{i=1}^{N} \frac{w_{t}^{i}}{\sum_{j=1}^{N} w_{t}^{j}} \delta_{x_{t}^{\prime}}\left(x_{t}\right) .
$$

## Particle filter - introductory example (I/II)

Consider a toy 1D localization problem.

Data


## Model

Dynamic model:

$$
x_{t+1}=x_{t}+u_{t}+v_{t}
$$

where $x_{t}$ denotes position, $u_{t}$ denotes velocity (known), $v_{t} \sim \mathcal{N}(0,5)$ denotes an unknown disturbance.

Measurements:

$$
y_{t}=h\left(x_{t}\right)+e_{t} .
$$

where $h(\cdot)$ denotes the world model (here the terrain height) and $e_{t} \sim \mathcal{N}(0,1)$ denotes an unknown disturbance.

Task: Find the state $x_{t}$ (position) based on the measurements $y_{1: t} \triangleq\left\{y_{1}, \ldots, y_{t}\right\}$ by computing the filter density $p\left(x_{t} \mid y_{1: t}\right)$.

## Particle filter - introductory example (II/II)




Highlights two key capabilities of the PF:

1. Automatically handles an unknown and dynamically changing number of hypotheses.
2. Work with nonlinear/nonGaussian models.

## Sequential Monte Carlo - particle filter



## SMC $=$ sequential importance sampling + resampling

1. Propagation: $x_{t}^{i} \sim p\left(x_{t} \mid x_{1: t-1}^{a_{t}^{i}}\right)$ and $x_{1: t}^{i}=\left\{x_{1: t-1}^{a_{t}^{i}}, x_{t}^{i}\right\}$.
2. Weighting: $\bar{w}_{t}^{i}=W_{t}\left(x_{t}^{i}\right)=p\left(y_{t} \mid x_{t}^{i}\right)$.
3. Resampling: $\mathbb{P}\left(a_{t}^{i}=j\right)=\bar{w}_{t-1}^{j} / \sum_{l} \bar{w}_{t-1}^{\prime}$.

The ancestor indices $\left\{a_{t}^{i}\right\}_{i=1}^{N}$ are very useful auxiliary variables! They make the stochasticity of the resampling step explicit.

## Application - indoor localization (I/III)

Aim: Compute the position of a person moving around indoors using sensors (inertial, magnetometer, radio) located in an ID badge and a map.


The sensors (IMU and radio) and the DSP are mounted inside an ID badge.

pdf for an office environment, the bright areas are rooms and corridors (i.e., walkable space).

## Application - indoor localization (II/III)



## Show movie

Johan Kihlberg, Simon Tegelid, Manon Kok and Thomas B. Schön. Map aided indoor positioning using particle filters. Reglermöte (Swedish Control Conference), Linköping, Sweden, June 2014.

## Application - indoor localization (III/III)

Aim: Compute the position using variations in the ambient magnetic field and the motion of the person (acceleration and angular velocities). All of this observed using sensors in a standard smartphone.


Movie - map making:
www. youtube.com/watch?v=enlMiUqPVJo
Movie - indoor positioning result

Arno Solin, Simo Särkkä, Juho Kannala and Esa Rahtu. Terrain navigation in the magnetic landscape: Particle filtering for indoor positioning. In Proceedings of the European Navigation Conference Helsinki, Finland, June, 2016.

Arno Solin, Manon Kok, Niklas Wahlström, Thomas B. Schön and Simo Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. IEEE Transactions on Robotics, 2018 (to appear).

## Particle MCMC = SMC + MCMC

A systematic way of combining SMC and MCMC.
Builds on an extended target construction.
Intuitively: SMC is used as a high-dimensional proposal mechanism on the space of state trajectories $\mathcal{X}^{T}$.

A bit more precise: Construct a Markov chain with $p\left(\theta, x_{1: T} \mid y_{1: T}\right)$ (or one of its marginals) as its stationary distribution. Also used for parameter learning.

## Exact approximations

Pioneered by the work:
Christophe Andrieu, Arnaud Doucet and Roman Holenstein, Particle Markov chain Monte Carlo methods, Journal of the Royal Statistical Society: Series B, 72:269-342, 2010.

## Outline

Aim: 1. Give a (hopefully) intuitive explanation of sequential Monte Carlo (SMC) for probabilistic modelling of dynamical systems. 2. Derive a new stochastic optimization method that can for example be used to learn unknown parameters in nonlinear SSMs.

1. Probabilistic modelling of dynamical systems
2. Sequential Monte Carlo (SMC)
3. Stochastic optimization
4. Some ongoing research snapshots (if there is time)

## Quasi-Newton - A non-standard take

Our problem is of the form (note change of notation...)

$$
\max _{x} f(x)
$$

Idea underlying (quasi-)Newton methods: Learn a local quadratic model $q\left(x_{k}, \delta\right)$ of the cost function $f(x)$ around the current iterate $x_{k}$

$$
q\left(x_{k}, \delta\right)=f\left(x_{k}\right)+g\left(x_{k}\right)^{\top} \delta+\frac{1}{2} \delta^{\top} H\left(x_{k}\right) \delta
$$

A second-order Taylor expansion around $x_{k}$, where

$$
\begin{aligned}
g\left(x_{k}\right) & =\left.\nabla f(x)\right|_{x=x_{k}}, \\
H\left(x_{k}\right) & =\left.\nabla^{2} f(x)\right|_{x=x_{k}}, \\
\delta & =x-x_{k} .
\end{aligned}
$$

## Available data

We have measurements of the

- cost function $f_{k}=f\left(x_{k}\right)$,
- and its gradient $g_{k}=g\left(x_{k}\right)$.

Question: How do we update the Hessian model?

Line segment connecting two adjacent iterates $x_{k}$ and $x_{k+1}$ :

$$
r_{k}(\tau)=x_{k}+\tau\left(x_{k+1}-x_{k}\right), \quad \tau \in[0,1] .
$$

## Useful basic facts

The fundamental theorem of calculus states that

$$
\int_{0}^{1} \frac{\partial}{\partial \tau} \nabla f\left(r_{k}(\tau)\right) \mathrm{d} \tau=\nabla f\left(r_{k}(1)\right)-\nabla f\left(r_{k}(0)\right)=\underbrace{\nabla f\left(x_{k+1}\right)}_{g_{k+1}}-\underbrace{\nabla f\left(x_{k}\right)}_{g_{k}}
$$

and the chain rule tells us that

$$
\frac{\partial}{\partial \tau} \nabla f\left(r_{k}(\tau)\right)=\nabla^{2} f\left(r_{k}(\tau)\right) \frac{\partial r_{k}(\tau)}{\partial \tau}=\nabla^{2} f\left(r_{k}(\tau)\right)\left(x_{k+1}-x_{k}\right)
$$

$$
\underbrace{g_{k+1}-g_{k}}_{=y_{k}}=\int_{0}^{1} \frac{\partial}{\partial \tau} \nabla f\left(r_{k}(\tau)\right) \mathrm{d} \tau=\int_{0}^{1} \nabla^{2} f\left(r_{k}(\tau)\right) \mathrm{d} \tau(\underbrace{x_{k+1}-x_{k}}_{s_{k}})
$$

## Result - the quasi-Newton integral

With the definitions $y_{k} \triangleq g_{k+1}-g_{k}$ and $s_{k} \triangleq x_{k+1}-x_{k}$ we have

$$
y_{k}=\int_{0}^{1} \nabla^{2} f\left(r_{k}(\tau)\right) \mathrm{d} \tau s_{k} .
$$

Interpretation: The difference between two consecutive gradients $\left(y_{k}\right)$ constitute a line integral observation of the Hessian.

Problem: Since the Hessian is unknown there is no functional form available for it.

## Solution 1 - recovering existing quasi-Newton algorithms

Existing quasi-Newton algorithms (e.g. BFGS, DFP, Broyden's method) assume the Hessian to be constant

$$
\nabla^{2} f\left(r_{k}(\tau)\right) \approx H_{k+1}, \quad \tau \in[0,1]
$$

implying the following approximation of the integral (secant condition)

$$
y_{k}=H_{k+1} s_{k} .
$$

Find $H_{k+1}$ by regularizing $H$ :

$$
\begin{aligned}
H_{k+1}=\min _{H} & \left\|H-H_{k}\right\|_{W}^{2}, \\
\text { s.t. } & H=H^{\top}, \quad H s_{k}=y_{k},
\end{aligned}
$$

Equivalently, the existing quasi-Newton methods can be interpreted as particular instances of Bayesian linear regression.

## Solution 2 - use a flexible nonlinear model

The approach used here is fundamentally different.
Recall that the problem is stochastic and nonlinear.
Hence, we need a model that can deal with such a problem.

Idea: Represent the Hessian using a Gaussian process learnt from data.

Two of the remaining challenges:

1. Can we use line integral observations when learning a GP?
2. How do we ensure that the resulting GP represents a Hessian?

## GP prior for the Hessian

## Stochastic quasi-Newton integral

$$
y_{k}=\int_{0}^{1} \underbrace{B\left(r_{k}(\tau)\right)}_{=\nabla^{2} f\left(r_{k}(\tau)\right)} s_{k} \mathrm{~d} \tau+e_{k},
$$

corresponds to noisy $\left(e_{k}\right)$ gradient observations.

Since $B(x) s_{k}$ is a column vector, the integrand is given by

$$
\operatorname{vec}\left(B(x) s_{k}\right)=\left(s_{k}^{\top} \otimes I\right) \operatorname{vec}(B(x))=\left(s_{k}^{\top} \otimes I\right) \operatorname{vec}(B(x))
$$

where vec $(B(x))=D \underbrace{\operatorname{vech}(B(x))}_{\widetilde{B}(x)}$.

Let us use a GP model for the unique elements of the Hessian

$$
\widetilde{B}(x) \sim \mathcal{G} \mathcal{P}\left(\mu(x), \kappa\left(x, x^{\prime}\right)\right)
$$

## Resulting stochastic qN integral and Hessian model

Summary: resulting stochastic quasi-Newton integral:

$$
y_{k}=\underbrace{\left(s_{k}^{\top} \otimes I\right) D}_{=\bar{D}_{k}} \int_{0}^{1} \widetilde{B}\left(r_{k}(\tau)\right) \mathrm{d} \tau+e_{k}
$$

with the following model for the Hessian

$$
\widetilde{B}(x) \sim \mathcal{G} \mathcal{P}\left(\mu(x), \kappa\left(x, x^{\prime}\right)\right)
$$

The Hessian can now be estimated using tailored GP regression.
Linear transformations (such as an integral or a derivative) of a GP results in a new GP.

## Resulting stochastic optimization algorithm

Standard non-convex numerical optimization loop with non-standard components.

Algorithm 1 Stochastic optimization

1. Initialization $(k=1)$
2. while not terminated do
(a) Compute a search direction $p_{k}$ using the current approximation of the gradient $g_{k}$ and Hessian $B_{k}$.
(b) Stochastic line search to find a step length $\alpha_{k}$ and set

$$
x_{k+1}=x_{k}+\alpha_{k} p_{k}
$$

(c) Set $k:=k+1$
(d) Update the Hessian estimate (tailored GP regression)

## 3. end while

## ex) Simple linear toy problem

Identify the parameters $\theta=(a, c, q, r)^{\top}$ in

$$
\begin{aligned}
x_{t+1} & =a x_{t}+w_{t}, & w_{t} & \sim \mathcal{N}(0, q), \\
y_{t} & =c x_{t}+e_{t}, & e_{t} & \sim \mathcal{N}(0, r) .
\end{aligned}
$$

Observations:

- The likelihood $L(\theta)=p\left(y_{1: T} \mid \theta\right)$ and its gradient $\nabla_{\theta} L(\theta)$ are available in closed form via standard Kalman filter equations.
- Standard gradient-based search algorithms applies.
- Deterministic optimization problem $\left(L(\theta), \nabla_{\theta} L(\theta)\right.$ noise-free $)$.


## ex) Simple linear toy problem



100 independent datasets.


Classical BFGS alg. for noisy observations of $L(\theta)$ and $\nabla L(\theta)$.

Clear blue - True system
Red - Mean value of estimate
Shaded blue - individual results
Red - Mean value of estimate
Shaded blue - individual results



GP-based BFGS alg. with noisy observations of $L(\theta)$ and $\nabla L(\theta) .30 / 41$

## ex) laser interferometry



The classic Michelson-Morley experiment from 1887.

Idea: Merge two light sources to create an interference pattern by superposition.

Two cases:

1. Mirror $B$ and $C$ at the same distance from mirror $A$.
2. Mirror $B$ and $C$ at different distances from mirror $A$.

## ex) laser interferometry

Dynamics: constant velocity model (with unknown force w)

$$
\binom{\dot{p}}{\dot{v}}=\left(\begin{array}{ll}
0 & 1 \\
0 & 0
\end{array}\right)\binom{p}{v}+\binom{0}{w} .
$$

Measurements: generated using two detectors

$$
\begin{array}{ll}
y_{1}=\alpha_{0}+\alpha_{1} \cos (\kappa p)+e_{1}, & e_{1} \sim \mathcal{N}\left(0, \sigma^{2}\right), \\
y_{2}=\beta_{0}+\beta_{1} \sin (\kappa p+\gamma)+e_{2}, & e_{2} \sim \mathcal{N}\left(0, \sigma^{2}\right) .
\end{array}
$$

Unknown parameters: $\theta=\left(\begin{array}{llllll}\alpha_{0} & \alpha_{0} & \beta_{0} & \beta_{1} & \gamma & \sigma\end{array}\right)^{\top}$.

Resulting maximum likelihood system identification problem

$$
\max _{\theta} p\left(y_{1: T} \mid \theta\right)
$$

## ex) laser interferometry




Research snapshots

## Snapshot 1 - scaling up to large problems

What is the key limitation of our GP-based optimization algorithm?

It does not scale to large-scale problems!

Still highly useful and competitive for small to medium sized problems.

We have developed a new technique that scales to very large problems.

## Snapshot 1 - scaling up to large problems

## Key innovations:

- Replace the GP with a matrix updated using fast Cholesky routines.
- Exploit a receding history of iterates and gradients akin to L-BFGS.
- An auxiliary variable Markov chain construction.


Training a deep CNN for MNIST data.


Logistic loss function with an L2 regularizer, gisette, 6000 observations and 5000 unknown variables.


Logistic loss function with an L2 regularizer, URL, 2396130 observations and 3231961 unknown variables.

Adrian Wills and Thomas B. Schön. Stochastic quasi-Newton with adaptive step lengths for large-scale problems. arXiv:1802.04310, February, 2018.

## Snapshot 2 - GP-based nonlinear state space model

"Inspired by the Gaussian process, enabled by the particle filter"

$$
\begin{aligned}
x_{t+1} & =f\left(x_{t}\right)+w_{t}, & \text { s.t. } & f(x) \sim \mathcal{G} \mathcal{P}\left(0, \kappa_{\eta, f}\left(x, x^{\prime}\right)\right), \\
y_{t} & =g\left(x_{t}\right)+e_{t}, & \text { s.t. } & g(x) \sim \mathcal{G} \mathcal{P}\left(0, \kappa_{\eta, g}\left(x, x^{\prime}\right)\right) .
\end{aligned}
$$

Results in a flexible non-parametric model where the GP prior takes on the role of a regularizer.

We can now find the posterior distribution

$$
p\left(f, g, Q, R, \eta \mid y_{1: T}\right),
$$

via some approximation (we use particle MCMC).

[^2]Andreas Svensson and Thomas B. Schön. A flexible state space model for learning nonlinear dynamical systems, Automatica,

## Snapshot 3 - The ASSEMBLE project and Birch

Aim: Automate probabilistic modeling of dynamical systems (and their surroundings) via a formally defined probabilistic modeling language.

Keep the model and the learning algorithms separated.
Create a market place for SMC-based learning algorithms (think CVX).
Birch — Our prototype probabilistic programming language.

[^3]
## Birch - our prototype probabilistic programming language

1. The basic idea of probabilistic programming is to equate probabilistic models with the programs that implement them.
2. Just as we can think of doing inference over models, we can think of doing inference over programs.

The particular PPL used here is Birch, which is currently being developed at Uppsala University.

Probabilistic and object-oriented language.
An early pre-release of Birch is available
birch-lang.org

## Snapshot 4 - The nonlinear SSM is just a special case...

Constructing an artificial sequence of intermediate target distributions for an SMC sampler is a powerful (quite possibly underutilized) idea.


Christian A. Naesseth, Fredrik Lindsten and Thomas B. Schön, Sequential Monte Carlo methods for graphical models. Advances in Neural Information Processing Systems (NIPS) 27, Montreal, Canada, December, 2014.

Fredrik Lindsten, Adam M. Johansen, Christian A. Naesseth, Bonnie Kirkpatrick, Thomas B. Schön, John Aston and Alexandre Bouchard-Côté. Divide-and-Conquer with Sequential Monte Carlo. Journal of Computational and Graphical Statistics (JCGS), 2017.

## Snapshot 5 - Spatio-temporal modelling

## Problem: predicting spatio-temporal

 processes with temporal patterns varying across spatial regions when data is obtained as a stream.A localized spatio-temporal covariance model.


The predictor can be updated sequentially with each new data point.

Muhammad Osama, Dave Zachariah and Thomas B. Schön. Learning localized spatio-temporal models from streaming data. In Proceedings of the 35th International Conference on Machine Learning (ICML), Stockholm, Sweden, July, 2018.

## Conclusion

Probabilistic modelling of nonlinear dynamical systems

$$
p\left(x_{0: T}, \theta, y_{1: T}\right)=\underbrace{\prod_{t=1}^{T} \underbrace{p\left(y_{t} \mid x_{t}, \theta\right)}_{\text {observation }}}_{\text {data distribution }} \underbrace{\prod_{t=1}^{T} \underbrace{p\left(x_{t} \mid x_{t-1}, \theta\right)}_{\text {dynamics }} \underbrace{p\left(x_{0} \mid \theta\right)}_{\text {state }} \underbrace{p(\theta)}_{\text {param. }}}_{\text {prior }}
$$

SMC provide approximate solutions to integration problems where there is a sequential structure present.

Stochastic optimization:

- Non-standard interpretation of quasi-Newton.
- Represent the Hessian using a Gaussian process.
- We can scale up to larg(er) problems.


[^0]:    Manon Kok, Jeroen D. Hol and Thomas B. Schön. Using inertial sensors for position and orientation estimation, Foundations and Trends of Signal Processing, 11(1-2):1-153, 2017.

[^1]:    Thomas B. Schön, Fredrik Lindsten, Johan Dahlin, Johan Wagberg, Christian A. Naesseth, Andreas Svensson and Liang Dai. Sequential Monte Carlo methods for system identification. In Proceedings of the 17th IFAC Symposium on System Identification (SYSID), Beijing, China, October 2015.

[^2]:    Frigola, Roger, Fredrik Lindsten, Thomas B. Schön, and Carl Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In Advances in Neural Information Processing Systems (NIPS), 2013.

[^3]:    Lawrence M. Murray, Daniel Lundén, Jan Kudlicka, David Broman and Thomas B. Schön. Delayed sampling and automatic
    Rao-Blackwellization of probabilistic programs. In Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS), Lanzarote, Spain, April, 2018.

