

6th Workshop on Sequential Monte Carlo Methods (SMC 2024)

Poster Sessions

Monday 13 May 2024 – Poster session 1	
Name:	Title:
Hany Abdulsamad	Nesting Particle Filters for Experimental Design in Dynamical Systems
Joshua Bon	Monte Carlo twisting for particle filters
Benjamin Boys	Tweedie Moment Projected Diffusions for Inverse Problems
John-Joseph Brady	Differentiable Particle Filtering Under Model Uncertainty:
Nicola Branchini	Generalizing self-normalized importance sampling
Matt Bright	Reversible Jump SMC
Mauro Camara Escudero	Integrator snippets for robust approximate manifold sampling
Abhishek Chakraborty	Dynamic Bayesian Real-time Kernel Density Estimation
Xiongjie Chen	Continuous-discrete Differentiable Particle Filters
Yuan Chen	TBC

Tuesday 14 May 2024 – Poster session 2	
Name:	Title:
Adrien Corenflos	Conditioning diffusion models by explicit forwards-backwards bridging
Sebastian Ertel	Continuous time Ensemble Kalman filters for infinite dimensional signals.
Omar Fabian González Hernandez	Enhancing Bayesian Filtering Through Iterative Nudging with the Ascendant gradient Method
Erdong Guo	TBC
Shu Huang	TBC
Yuga Iguchi	Parameter Estimation with Increased Precision for Elliptic and Hypo-elliptic Diffusions
Minas Karamanis	Persistent Sequential Monte Carlo
Joona Karjalainen	On the Forgetting of Particle Filters
Minhye Kim	Dynamic Integration of Compartmental Modeling and Assimilation Techniques for Seasonal Influenza Forecasting
Jiaxi Li	Learning Differentiable Particle Filter on the Fly

Wednesday 15 May 2024 – Poster session 3	
Name:	Title:
Pinak Mandal	Numerical explorations of nonlinear filter stability using Sinkhorn divergence
Nicolò Margaritella	Estimating the local True Discovery Rate and its uncertainty with ABC SMC
Hugo Marival	The Importance Markov Chain
Alessandro Mastrototaro	Online Variational Sequential Monte Carlo
Thorben Pieper	Simulation of Infinite Dimensional Diffusion Bridges

Hans Reimann	The Diffusion Score Matching Kalman Filter: A Cheap and Provably Robust Adaptation for Observation Noise Mis-Specification
Man Ho Suen	Linearization approach for aggregated landslides data
Jia Le Tan	Pareto-Smoothed Sequential Monte Carlo
Zheng Zhao	On Feynman–Kac training of partial Bayesian neural networks

NESTING PARTICLE FILTERS FOR EXPERIMENTAL DESIGN IN DYNAMICAL SYSTEMS

HANY ABDULSAMAD

We propose a novel approach to sequential Bayesian Experimental Design (BED) for non-exchangeable data that formulates it as risk-sensitive policy optimization. We develop the Inside-Out SMC² algorithm that uses a nested sequential Monte Carlo (SMC) estimator of the expected information gain and embeds it into a particle Markov chain Monte Carlo (pMCMC) framework to perform gradient-based policy optimization. This is in contrast to recent approaches that rely on biased estimators of the expected information gain (EIG) to amortize the cost of experiments by learning a design policy in advance. Numerical validation on a set of dynamical systems showcases the efficacy of our method in comparison to other state-of-the-art strategies.

MONTE CARLO TWISTING FOR PARTICLE FILTERS

JOSHUA BON

We consider the problem of designing efficient particle filters for twisted Feynman--Kac models. Particle filters using twisted models can deliver low error approximations of statistical quantities and such twisting functions can be learnt iteratively. Practical implementations of these algorithms are complicated by the need to (i) sample from the twisted transition dynamics, and (ii) calculate the twisted potential functions. We expand the class of applicable models using rejection sampling for (i) and unbiased approximations for (ii) using a random weight particle filter. We characterise the average acceptance rates within the particle filter in order to control the computational cost, and analyse the asymptotic variance. Empirical results show the mean squared error of the normalising constant estimate in our method is smaller than a memory-equivalent particle filter but not a computation-equivalent filter. Both comparisons are improved when more efficient sampling is possible which we demonstrate on a stochastic volatility model.

TWEEDIE MOMENT PROJECTED DIFFUSIONS FOR INVERSE PROBLEMS

BENJAMIN BOYS

Diffusion generative models unlock new possibilities for inverse problems as they allow for the incorporation of strong empirical priors into the process of scientific inference. Recently, diffusion models are repurposed for solving inverse problems using Gaussian approximations to conditional densities of the reverse process via Tweedie's formula to parameterise the mean, complemented with various heuristics. To address various challenges arising from these approximations, we leverage higher order information using Tweedie's formula and obtain a statistically principled approximation. We further provide a theoretical guarantee specifically for posterior sampling which can lead to better theoretical understanding of diffusion-based conditional sampling. Finally, we illustrate the empirical effectiveness of our approach for general linear inverse problems on toy synthetic examples as well as image restoration. We show that our method (i) removes any time-dependent step-size hyperparameters required by earlier methods, (ii) brings stability and better sample quality across multiple noise levels, (iii) is the only method that works in a stable way with variance exploding (VE) forward processes as opposed to earlier works.

DIFFERENTIABLE PARTICLE FILTERING UNDER MODEL UNCERTAINTY

JOHN-JOSEPH BRADY

Differentiable particle filters are an emerging class of models that combine sequential Monte Carlo techniques with the flexibility of neural networks to perform state space inference. This paper concerns the case where the system may switch between a bank of several discretely labeled regimes. Currently there are no algorithms in the literature that effectively learn both the dynamics of individual models and the switching process simultaneously. In this paper, we propose the regime learning differentiable particle filter, a novel algorithm, to solve this problem. On a pair of simulated data experiments, we demonstrate competitive performance compared to the previous state-of-the-art algorithms.

GENERALIZING SELF-NORMALIZED IMPORTANCE SAMPLING

NICOLA BRANCHINI

TBC

REVERSIBLE JUMP SMC

MATT BRIGHT

The problem of approximating or sampling from an unknown distribution is compounded when that distribution is governed by a mixture of discrete and continuous parameters. The method of Reversible Jump MCMC [1], is a commonly used tool in such situations. We propose to counter the slow convergence of RJMCMC by adapting the algorithm to a Sequential Monte Carlo framework.

We compare three approaches:

- Updating the weights of particles proposed by RJMCMC at each iteration using the acceptance ratio computed under RJMCMC.
- Using particles proposed by RJMCMC only if they are accepted, otherwise passing the same particle on to the next iteration. Weights are not updated (effectively running parallel RJMCMC chains)
- Including an accept/reject step as above, updating weights using the acceptance ratio only for accepted particles.

We investigate the behaviour of these three approach for the algorithm developed by Richardson and Green for Univariate Gaussian Mixture Models [2], which suggests that inclusion of an explicit accept/reject step reduces the variance of estimators and the need for resampling at each iteration, at the expense of a slower apparent rate of convergence.

INTEGRATOR SNIPPETS FOR ROBUST APPROXIMATE MANIFOLD SAMPLING

MAURO CAMARA ESCUDERO

Efficient sampling from probability distributions concentrated around a lower-dimensional manifold is crucial in numerous applications arising in machine learning and statistical physics. We propose Gibbs Hug Markov Snippet (GHUMS) a robust and parallelisable inference algorithm, based on recent work on Integrator Snippets, that is well-adapted for these problems and show its superiority in terms of efficiency and robustness.

DYNAMIC BAYESIAN REAL-TIME KERNEL DENSITY ESTIMATION

ABHISHEK CHAKRABORTY

I will be presenting a working paper, which hopefully will be submitted by the time the workshop and the summer school will be held, titled "Dynamic Bayesian Real-time Kernel Density Estimation", in which we are using Bayesian Filtering and some SMC methods for capturing the variability in kernel density for streaming batches of data. We have experimented with Kalman Filters, Particle Filters and Extended Kalman Filters, all of which have somewhat performed better in terms of achieving better prediction accuracy for density estimation compared to Temporal Adaptive Kernel Density Estimation.

CONTINUOUS-DISCRETE DIFFERENTIABLE PARTICLE FILTERS

XIONGJIE CHEN

Differentiable particle filters are a recently emerging family of data-adaptive sequential Monte Carlo methods built with neural networks and optimised by gradient descent. One limitation of the existing differentiable particle filters is that they can only infer latent states at predefined discrete (often evenly-spaced) time instances. In our work, we extend differentiable particle filters to continuous-discrete settings, where latent states evolve in continuous time and observations are emitted at discrete time steps. The proposed continuous-discrete differentiable particle filters (CD-DPFs) are constructed with expressive machine learning tools such as neural stochastic differential equations and normalising flows. We validate the effectiveness of the CD-DPF and compare its performance with state-of-the-art baseline methods on benchmark continuous-discrete filtering and observation prediction/interpolation datasets.

TBC

YUAN CHEN

State-space models (SSMs) are widely used in statistics, econometrics, and signal processing. In most applications, the model depends on unknown static parameters that need to be inferred from the data. An idea of performing parameter estimation in SSMs involves extending the model to include its parameters by treating them as additional state variables. We study the sequence of extended joint filtering distributions and provide convergence guarantees to concentrate towards target values. We use the standard particle filter for online predictions and online inference in SSMs, and have shown that our online optimal predictions of state variables are in some sense to the optimal filtering distribution. As a byproduct, we generalized our results to iterative filtering.

CONDITIONING DIFFUSION MODELS BY EXPLICIT FORWARDS- BACKWARDS BRIDGING

ADRIEN CORENFLOS

Diffusion models have emerged as powerful generative procedures in many machine learning and statistical applications. At heart, they rely on learning a forward-noising backward-denoising pair of stochastic differential equations, transporting the target distribution of interest to and back from a reference measure. Given an unconditional diffusion model, using them to perform conditional simulation is still largely an open question and is typically achieved by learning conditioning drifts to the backward SDE after the fact.

In this work, we leverage the dynamical structure of diffusion models to show that conditioning therein can be written as a partial SDE bridge, constructed by running a preliminary forward simulation, and then a backward particle filter corresponding to the bridging.

To improve the performance of the method, we embed this forward-backwards bridging procedure within a conditional sequential Monte Carlo (cSMC) Markov chain, thereby recovering an "approximately exact" sampling procedure for the conditioned diffusion model.

We will showcase preliminary results, demonstrating the promising nature of the method.

CONTINUOUS TIME ENSEMBLE KALMAN FILTERS FOR INFINITE DIMENSIONAL SIGNALS

SEBASTIAN ERTEL

Continuous time filtering problems where the signal is determined by an SPDE, but observations are finite dimensional are considered. We derive a mean field Ensemble Kalman–Bucy filter (EnKBF) suitable to this setting by a constant gain approximation of the corresponding infinite dimensional Feedback Particle Filter.

Well posedness of the (mean-field) EnKBF is proven by Picard iterations in the observation function. Finally an (almost) optimal propagation of chaos result is obtained for the EnKBF.

ENHANCING BAYESIAN FILTERING THROUGH ITERATIVE NUDGING WITH THE ASCENDANT GRADIENT METHOD

OMAR FABIAN GONZÁLEZ HERNANDEZ

The ascendant method, implemented through the nudging technique, proves to be a valuable tool in enhancing the likelihood function within a Bayesian filtering framework. By iteratively adjusting model parameters toward the observed data, nudging facilitates a more accurate alignment between model predictions and actual measurements. This iterative refinement, guided by the ascendant method, not only enhances the consistency of the filtering process but also effectively increases the likelihood of the observed data given the model. In doing so, the ascendant method via nudging contributes to a more robust and reliable Bayesian filtering framework, aligning model outputs more closely with real-world observations.

TBC

ERDONG GUO

TBC

TBC

SHU HUANG

We introduce an inferential framework for a wide class of partially observed stochastic differential equations (SDEs), with drift satisfying a global one-sided Lipschitz condition with at most polynomial growth. Past work has witnessed numerical schemes that employ explicit pseudo-likelihoods which preserve critical properties of the SDE. We formalise the computation of the implied partial pseudo-likelihood as a Feynman-Kac path model, allowing for its efficient estimation by the controlled Sequential Monte Carlo algorithm. In the Bayesian framework, we introduce a partial-pseudo-likelihood-driven Markov Chain Monte Carlo method that targets the posterior distribution. We propose to numerically maximize the estimated partial pseudo-likelihood to give point estimates that enjoys good asymptotic properties and at low computational cost. We include numerical examples that demonstrate our proposed methodology for different numerical schemes on both hypoelliptic and elliptic diffusions.

PARAMETER ESTIMATION WITH INCREASED PRECISION FOR ELLIPTIC AND HYPO-ELLIPTIC DIFFUSIONS

YUGA IGUCHI

This work aims at making a comprehensive contribution in the general area of parametric inference for discretely observed diffusion processes. Established approaches for likelihood-based estimation invoke a time-discretisation scheme for the approximation of the intractable transition dynamics of the Stochastic Differential Equation (SDE) model over finite time periods. The scheme is applied for a step-size that is either user-selected or determined by the data. Recent research has highlighted the critical effect of the choice of numerical scheme on the behaviour of derived parameter estimates in the setting of hypo-elliptic SDEs. In brief, in our work, first, we develop two weak second order sampling schemes (to cover both hypo-elliptic and elliptic SDEs) and produce a small time expansion for the density of the schemes to form a proxy for the true intractable SDE transition density. Then, we establish a collection of analytic results for likelihood-based parameter estimates obtained via the formed proxies, thus providing a theoretical framework that showcases advantages from the use of the developed methodology for SDE calibration. We present numerical results from carrying out classical or Bayesian inference, for both elliptic and hypo-elliptic SDEs.

PERSISTENT SEQUENTIAL MONTE CARLO

MINAS KARAMANIS

Sequential Monte Carlo (SMC) methods are popular tools for Bayesian inference, particularly in complex, non-linear settings. However, their computational cost and high-variance estimates limit their scalability. We introduce Persistent Sequential Monte Carlo (PSMC), a novel extension that overcomes these limitations by leveraging information from all past iterations. PSMC achieves this through an inclusive reweighting and resampling strategy, leading to a richer sample set, lower variance estimates for the marginal likelihood, and faster diversification of particles. By maintaining diverse pool of persistent particles, PSMC effectively mitigates particle impoverishment and variance inflation. Rigorous numerical experiments and real-world applications demonstrate the significant advantages of PSMC, offering a more robust and efficient sampling process for challenging Bayesian inference problems.

ON THE FORGETTING OF PARTICLE FILTERS

JOONA KARJALAINEN

We study the forgetting properties of the particle filter when its state - the collection of particles - is regarded as a Markov chain. Under a strong mixing assumption on the particle filter's underlying Feynman-Kac model, we find that the particle filter is exponentially mixing, and forgets its initial state in $O(\log N)$ 'time', where N is the number of particles and time refers to the number of particle filter algorithm steps, each comprising a selection (or resampling) and mutation (or prediction) operation. We present an example which suggests that this rate is optimal. In contrast to our result, available results to-date are extremely conservative, suggesting $O(\alpha N)$ time steps are needed, for some $\alpha > 1$, for the particle filter to forget its initialisation. We also study the conditional particle filter (CPF) and extend our forgetting result to this context. We establish a similar conclusion, namely, CPF is exponentially mixing and forgets its initial state in $O(\log N)$ time. To support this analysis, we establish new time-uniform L_p error estimates for CPF, which can be of independent interest.

DYNAMIC INTEGRATION OF COMPARTMENTAL MODELING AND ASSIMILATION TECHNIQUES FOR SEASONAL INFLUENZA FORECASTING

MINHYE KIM

This research presents an innovative method that merges two distinct data assimilation strategies for forecasting seasonal influenza in real-time, leveraging the SIR model's framework. By analyzing Influenza-Like-Illness (ILI) data from Korea between 2016 and 2021, our method demonstrated its capability to accurately forecast flu incidences for the forthcoming week, offering significant implications for enhancing public health decisions and managing medical resources efficiently during epidemic periods.

LEARNING DIFFERENTIABLE PARTICLE FILTER ON THE FLY

JLAXI LI

Differentiable particle filters are an emerging class of sequential Bayesian inference techniques that use neural networks to construct components in state space models. Existing approaches are mostly based on offline supervised training strategies. This leads to the delay of the model deployment and the obtained filters are susceptible to distribution shift of test-time data. In this paper, we propose an online learning framework for differentiable particle filters so that model parameters can be updated as data arrive. The technical constraint is that there is no known ground truth state information in the online inference setting. We address this by adopting an unsupervised loss to construct the online model updating procedure, which involves a sequence of filtering operations for online maximum likelihood-based parameter estimation. We empirically evaluate the effectiveness of the proposed method, and compare it with supervised learning methods in simulation settings including a multivariate linear Gaussian state-space model and a simulated object tracking experiment.

NUMERICAL EXPLORATIONS OF NONLINEAR FILTER STABILITY USING SINKHORN DIVERGENCE

PINAK MANDAL

Particle filters and ensemble Kalman filters are widely used in data assimilation but in the case of deterministic systems, only a few theoretical results for their stability are available. In this study, we explore the distance between filtering distributions starting from different initial distributions as a function of time using Wasserstein metric. We investigate the exponential nature and robustness of filter stability by varying two crucial parameters of the nonlinear filtering problem, namely the observation gap and the observation noise. We also establish numerically a relation between filter stability and filter RMSE.

ESTIMATING THE LOCAL TRUE DISCOVERY RATE AND ITS UNCERTAINTY WITH ABC SMC

NICOLÒ MARGARITELLA

We developed a "less empirical" Bayes method for estimating the local true discovery rate (lTDR) for large scale inference. The new approach is based on ABC SMC which allows us to approximate the posterior distribution of the lTDR and of other quantities of interests (e.g. the probability of H_0 and the density function of H_1). The method achieves optimal classification results under both sparsity and density scenarios and different effect sizes. The uncertainty evaluation of the true and false discoveries that ABC lFDR provides is a promising feature for the analysis of complex high dimensional scenarios such as group networks which arise in many fields (e.g. neuroscience, genetics).

THE IMPORTANCE MARKOV CHAIN

HUGO MARIVAL

The Importance Markov chain is a novel algorithm bridging the gap between rejection sampling and importance sampling, moving from one to the other through a tuning parameter. Based on a modified sample of an instrumental Markov chain targeting an instrumental distribution (typically via a MCMC kernel), the Importance Markov chain produces an extended Markov chain where the marginal distribution of the first component converges to the target distribution. For example, when targeting a multimodal distribution, the instrumental distribution can be chosen as a tempered version of the target which allows the algorithm to explore its modes more efficiently. We obtain a Law of Large Numbers and a Central Limit Theorem as well as geometric ergodicity for this extended kernel under mild assumptions on the instrumental kernel. Computationally, the algorithm is easy to implement and preexisting libraries can be used to sample from the instrumental distribution.

ONLINE VARIATIONAL SEQUENTIAL MONTE CARLO

ALESSANDRO MASTROTOTARO

In this work, we build upon the variational sequential Monte Carlo (VSMC) method, which provides computationally efficient and accurate model parameter estimation and Bayesian latent-state inference by combining particle methods and variational inference. While standard VSMC operates in the offline mode, by re-processing repeatedly a given batch of data, we distribute the approximation of the gradient of the VSMC surrogate ELBO in time using stochastic approximation, allowing for online learning in the presence of streams of data. This results in an algorithm, online VSMC, that is capable of performing efficiently, entirely on-the-fly, both parameter estimation and particle proposal adaptation. In addition, we provide rigorous theoretical results describing the algorithm's convergence properties as the number of data tends to infinity as well as numerical illustrations of its excellent convergence properties and usefulness also in batch-processing settings.

Joint work with Jimmy Olsson (KTH)

SIMULATION OF INFINITE DIMENSIONAL DIFFUSION BRIDGES

THORBEN PIEPER

Consider a mild solution to a stochastic partial differential equation (SPDE), conditioned on hitting an arbitrary endpoint at some fixed point in time. We call the arising process an 'infinite dimensional diffusion bridge' in accordance to the finite dimensional case of a conditioned stochastic differential equation.

In this work, we first derive an expression of the law of an infinite dimensional diffusion bridge. The derivation is based on solving martingale problems corresponding to infinite dimensional Kolmogorov operators and an application of Doob's h-transform.

Secondly, we introduce a way of simulating infinite dimensional diffusion bridges. It is well known that even in the finite dimensional case, these processes are intractable, rendering a direct simulation unattainable.

We therefore propose a Monte Carlo method in which samples are drawn from a proposal distribution corresponding to a tractable infinite dimensional diffusion process, which we term a 'guided process'. We show that the likelihood ratio between the diffusion bridge and guided process is well defined, thus enabling simulation by a simple accept-reject type algorithm.

The performance of the proposed solution will be presented in numerical experiments.

This poster is based on joint work with Frank van der Meulen (VU Amsterdam) and Aad van der Vaart (TU Delft).

THE DIFFUSION SCORE MATCHING KALMAN FILTER: A CHEAP AND PROVABLY ROBUST ADAPTATION FOR OBSERVATION NOISE MIS-SPECIFICATION

HANS REIMANN

Bayesian posterior distributions have highly desirable properties in Zellner's information optimality and the Bernstein-von-Mises theorem assuming a model is well-specified. However, as soon as there is reason for concern about mis-specification of the Bayesian model, these results may also no longer hold. Especially regarding mis-specification in heavy tail behaviour of the observation likelihood, so in potential observation outliers, Kullback-Leibler divergence at the heart of Bayes theorem is volatile. Yet, estimating observation noise distributions is a frequent challenge in data assimilation often subject to strong assumptions for feasibility of the problem.

With recent advances on scalable, robust generalized posteriors via diffusion score matching, new opportunities for the Bayesian inverse inference problem in the analysis step of the Kalman filter arose. We derived an adapted closed form analysis step of the popular Kalman filter based on these recent results with provable robustness to outliers in the framework of Huber's epsilon-contaminated. Next to theoretical results and derivations for the novel algorithm, we provide experimental results to showcase impact and potential.

LINEARIZATION APPROACH FOR AGGREGATED LANDSLIDES DATA

MAN HO SUEN

In spatial statistics, it is not uncommon to have spatial misalignment in observed responses at point locations and covariates data at various resolutions and shapes. One of the common approaches is to aggregate the point observations into count data with respect to the area polygon. One of the popular approaches in landslide literature is to aggregate based on slope units that cluster landslide observations beneath the surface. This takes away the point location information and introduces both bias and uncertainty. Starting with a Poisson point process, the domain is discretised into subspaces. The definition of these subspaces can be flexible based on various scenarios. Assuming the intensity of the process is log-linear, an implementation trick is used and the first-order Taylor linearization in the INLA and inlabru R packages. The approximation bias is computed with the help of the omitted second-order terms. This turns out to provide insights into improving the modelling of aggregated data.

PARETO-SMOOTHED SEQUENTIAL MONTE CARLO

JIA LE TAN

Importance Sampling (IS) is recognised for its theoretically unbiased nature. However, challenges arise when the proposal distribution markedly deviates from the target distribution, primarily due to the high variance in the weights it generates. In scenarios where a sampled point from a low-likelihood region in the proposal distribution aligns with a high-likelihood in the target distribution, there is a disproportionate 'weight stealing' effect, leading to a reduced effective sample size. This issue persists in higher dimensions, even when the proposal and target distributions are relatively similar. To address this, Vehtari et al. (2022) introduced Pareto-Smoothing (PS) within IS, culminating in the development of Pareto-Smoothed Importance Sampling (PSIS). This technique has been effective in substantially reducing weight variance with only a minimal increase in bias, especially when compared to methods like Truncated Importance Sampling (TIS).

Building on the foundation laid by Vehtari et al., our research integrates the PS methodology within the Sequential Monte Carlo (SMC) framework. Our goal is to replicate the enhancements observed in IS for SMC. In particular, we selectively apply Pareto-Smoothing to overcome challenges to SMC identical to those of IS aforementioned, such as sample impoverishment and weight degeneracy.

A pivotal aspect of our study is the investigation of diverse methods for embedding Pareto-Smoothing into the SMC framework. We showcase various integration strategies and evaluate the efficacy of these novel methods using environmental models. Through this comparative analysis, we aim to pinpoint the most effective Pareto-Smoothing techniques for optimising SMC, specifically tailored to varied environmental modelling contexts.

ON FEYNMAN–KAC TRAINING OF PARTIAL BAYESIAN NEURAL NETWORKS

ZHENG ZHAO

Recently, partial Bayesian neural networks (pBNNs), which only consider a subset of the parameters to be stochastic, were shown to perform competitively with full Bayesian neural networks. However, pBNNs are often multi-modal in the latent-variable space and thus challenging to approximate with parametric models. To address this problem, we propose an efficient sampling-based training strategy, wherein the training of a pBNN is formulated as simulating a Feynman--Kac model. We then describe variations of sequential Monte Carlo samplers that allow us to simultaneously estimate the parameters and the latent posterior distribution of this model at a tractable computational cost. We show on various synthetic and real-world datasets that our proposed training scheme outperforms the state of the art in terms of predictive performance.