MONTE CARLO METHODS FOR BAYESIAN PARAMETER INFERENCE

KASPER BÅGMARK AND RUBEN SEYER

1. Project description

Many time-series models contain a latent state observed through noisy measurements, together with unknown model parameters. Recent work uses deep neural networks to approximate the predictive state distributions and a differentiable surrogate for the parameter likelihood [1, 2].

At observation time k, the relevant likelihood contribution is

$$\mathcal{Z}_k(\vartheta) = p_\vartheta(o_k \mid o_{1:k-1}) = \mathbb{E}_{X \sim p^-_{k,\vartheta}}\big[\mathcal{L}_\vartheta(o_k \mid X)\big],$$

where $p^-_{k,\vartheta}$ is the predictive distribution of the latent state. The current method estimates these expectations by importance sampling and uses the resulting values to train a neural likelihood surrogate.

This thesis will investigate Markov chain Monte Carlo (MCMC), primarily Hamiltonian Monte Carlo (HMC), as an alternative. HMC can sample from the differentiable, potentially unnormalized distributions represented by the learned filter. The methods will be compared in terms of accuracy, robustness, and computational cost, and through their effect on the learned likelihood and parameter posterior. Trained networks and an existing inference pipeline will be provided, so the main focus is Monte Carlo sampling, computational statistics, and machine learning.

1.1. Main tasks.

The core project is to:

(1) study likelihood-based Bayesian inference, importance sampling, MCMC, and HMC;
(2) implement HMC for the learned predictive distributions and assess its convergence and sampling efficiency;
(3) benchmark the likelihood estimates using a linear Gaussian model and a nonlinear low-dimensional model with accurate references; and
(4) retrain the likelihood surrogate and evaluate the resulting parameter estimates and posterior distribution.

1.2. Possible extensions.

One extension may be pursued, such as tempered or bridge sampling, uncertainty-aware surrogate training, propagation of likelihood error to the parameter posterior, or comparison with another MCMC method.

1.3. Student profile and supervision.

The student should have a solid background in probability and statistics and be comfortable with numerical programming. Courses in Bayesian or computational statistics, machine learning, optimization, or Monte Carlo methods are particularly relevant. Experience with programming is expected. No prior knowledge of stochastic analysis, PDEs, or BSDEs is required.

The project will be carried out in collaboration with Vrije Universiteit Amsterdam. It combines a literature study with implementation and numerical evaluation in close collaboration with the supervisors. The project will be supervised by Kasper Bågmark at Vrije Universiteit Amsterdam and Ruben Seyer at the Department of Mathematical Sciences, Chalmers University of Technology and the University of Gothenburg.

References

[1] K. Bågmark, A. Andersson, and S. Larsson, *Nonlinear filtering based on density approximation and deep BSDE prediction*, arXiv:2508.10630, 2025.

[2] K. Bågmark and F. Rydin, *Neural likelihood surrogates for parameter inference via log-density PDE*, ICLR 2026 Workshop on AI and Partial Differential Equations, 2026. OpenReview.

[3] A. Gelman and X.-L. Meng, *Simulating normalizing constants: From importance sampling to bridge sampling to path sampling*, *Statistical Science*, 13(2):163–185, 1998.

[4] R. M. Neal, *MCMC using Hamiltonian dynamics*, in S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng (eds.), *Handbook of Markov Chain Monte Carlo*, Chapman and Hall/CRC, 2011.

Email address: k.bagmark@vu.nl

URL: https://bagmark.github.io

Email address: rubense@chalmers.se

URL: https://ruben.seyer.se